National Research
Council Canada

Conseil national
de recherches Canada

Associate Committee
on Artificial Intelligence

Comité associé de
l'intelligence artificielle

Artificial Intelligence in Canada

A Description by Members of the ACAI

January 1991

Canada

NRC No. 31789

91-05447/

Copyright 1991 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables from this report, provided that the source of such material is fully acknowledged.

Copyright 1991 par
Conseil national de recherches du Canada

Il est permis de citer de courts extraits et de reproduire des figures ou tableaux du présent rapport, à condition d'en identifier clairement la source.

Additional copies are available free of charge from:

Editorial Office, Room 301
Institute for Information Technology
National Research Council of Canada
Ottawa, Ontario, Canada
K1A 0R6

Des exemplaires supplémentaires gratuits peuvent être obtenus à l'adresse suivante :

Bureau des publications, Pièce 301
Institut de technologie de l'information
Conseil national de recherches du Canada
Ottawa (Ontario) Canada
K1A 0R6

Disponible en français, CNRC n° 31790.

Artificial Intelligence in Canada

A Description by Members of the ACAI

January 1991

Table of Contents

Preface v

Section 1: Executive Summary 1
  Executive Summary 3

Section 2: Introduction 15
  Introduction 17

Section 3: Consortia, Collaborative and Funding Organizations 21
  Canadian Institute for Advanced Research 23
  Centre de recherche informatique de Montréal 31
  Information Technology Research Centre 34
  PRECARN 38
  PRECARN'S Institute for Robotics and Intelligent Systems Network 43
  Government Support Programs for AI R&D in Canadian Industry 51
  Natural Sciences and Engineering Research Council 56

Section 4: AI in Universities, Government Departments, and Agencies 59
  Canadian Space Agency 61
  Department of Communications, Canadian Workplace Automation Research Centre 66
  Department of Communications, Communications Research Centre 69
  Department of National Defence 70
  Energy, Mines and Resources Canada 72
  Government Consulting Group 84
  National Research Council 85
  Provincial Research Organizations 91
  Transport Canada 96
  University Research in AI: An Overview 99
  Research in the Knowledge Sciences at the University of Calgary 102
  Recent AI Activities at the University of Guelph 110
  Artificial Intelligence at Queen's 116
  University of Saskatchewan 122
  Simon Fraser University 124
  Research in Biological and Computational Vision at the University of Toronto 128

Section 5: AI in Business and Industry 137
  Acquired Intelligence Inc 139
  Applied AI Systems, Inc 142
  Bell-Northern Research 144
  CAE Electronics Ltd. 147
  Canadian Pacific 151
  CompEngServ Ltd. 153
  Digital Equipment Corporation 155
  Hydro-Québec 157
  Artificial Intelligence Work at MacDonald Dettwiler and Associates Ltd. 165
  New Brunswick Power Commission 168
  Price Waterhouse 171
  Royal Bank of Canada 173
  Shell Canada 174

Section 6: Appendices 177
  Appendix 1: Canadian AI Capability Survey: December 1989 179
  Appendix 2: Canadian Society for Computational Studies of Intelligence .. 192
  Appendix 3: Computational Intelligence 193
  Appendix 4: National Research Council, Associate Committee on Artificial Intelligence 194
  Appendix 5: References and Related Reports 197

Section 7: Index 199

Preface

Associate Committees established by the National Research Council Canada have been used for over fifty years on many topics to act as a focal point and forum in Canada for issues of national concern to the engineering and scientific communities.

Artificial Intelligence (AI) is a subject of current and worldwide interest and many countries have strong R&D programs in AI. Accordingly, the National Research Council (NRC) Associate Committee on Artificial Intelligence (ACAI) was established in April 1987 to act as a forum for discussion and planning in this area at the national level.

The intent of this report is to provide a composite description of Artificial Intelligence research, development, and application as it exists throughout Canada today, based on knowledge and information readily available to members of the ACAI. While the report contains many descriptions and examples of AI work and capability, it is apparent at this same time that no report can be fully representative. In producing this report it is the committee's expectation that it will be useful as a compendium that will provide guidance and encouragement to others in the substantial and growing community of AI researchers and practitioners.

Copies of the committee's first and previous report, *The Social Context of Artificial Intelligence*, are also available in both French and English from the address listed on the inside front cover.

The ACAI wishes to acknowledge and express appreciation to the considerable number of organizations who have contributed information for inclusion in the report. The assistance of the Editorial Office, Institute for Information Technology (formerly the Division of Electrical Engineering), NRC, in the editing and final preparation of the ACAI reports is also appreciated and gratefully acknowledged.

As a perhaps interesting comment, this report has been almost entirely produced using optical character recognition (OCR) (as a universal interface) to readily enter the many contributions received into a modern desktop publishing system for conversion to a common style and format.

To facilitate its use as a reference, an extensive index is provided at the end of the report.

Ottawa, Ontario
July 1990

Section 1

Executive Summary

Executive Summary

For convenience the descriptions in the report are divided into three main sections:

- Consortia and other organizations for collaborative research and development in artificial intelligence.
- Research and development activities in universities, government departments, agencies, and provincial research councils, including those for application development and technology transfer.
- Application development and use in business and industry.

Figure 1 provides a pictorial representation of the report content by reference to the many companies and research organizations that have generously contributed the information provided. In addition, a detailed index is provided at the end of the report to assist the reader wishing to obtain information on any specific topic, for which references may often be found in more than one section of the report. Appendix 1 provides results on a survey of Canadian AI capability. This indicates that there are more than 1000 full-time and several hundred part-time individuals in Canada involved in AI research, development, and application in more than 300 organizations across Canada.

Consortia, Collaborative, and Funding Organizations

Canadian Institute for Advanced Research

The AI and Robotics (AIR) program of the Canadian Institute for Advanced Research (CIAR) was established in 1984. This recognized AI as an emerging field of considerable importance to business and industry, and recognized the need for supplementary funding in order to establish research excellence and critical mass for AI and robotics research in Canadian universities.

Through its support, the 15 Fellows of the AIR Program and others have contributed to all aspects of the study of intelligent systems. World-class centers of high repute have been established at three major centers located at the University of Toronto, McGill University, and the University of British Columbia, with additional centers of strength at The University of Western Ontario, Queen's University, and the University of Alberta.

Centre de recherche informatique de Montréal

The Centre de recherche informatique de Montréal (CRIM) is a computer research center

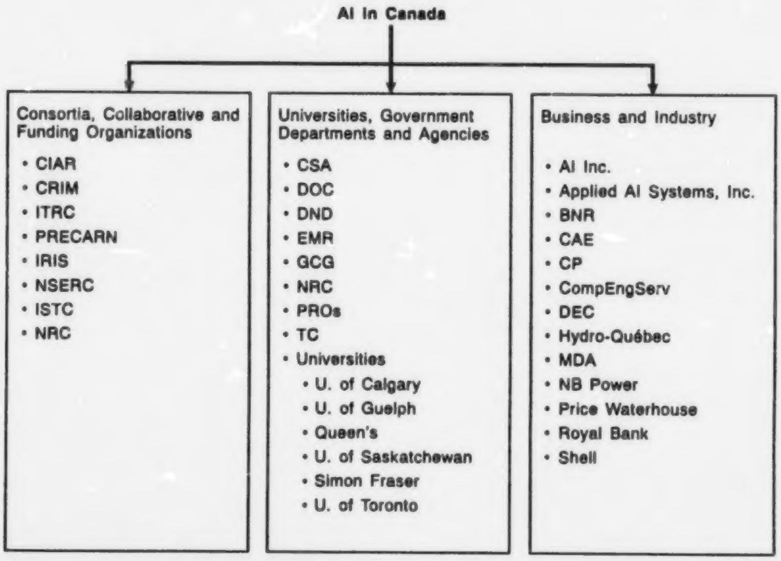
AI in Canada
Consortia, Collaborative and Funding Organizations
• CIAR
• CRIM
• ITRC
• PRECARN
• IRIS
• NSERC
• ISTC
• NRC
Universities, Government Departments and Agencies
• CSA
• DOC
• DND
• EMR
• GCG
• NRC
• PROs
• TC
• Universities
• U. of Calgary
• U. of Guelph
• Queen's
• U. of Saskatchewan
• Simon Fraser
• U. of Toronto
Business and Industry
• AI Inc.
• Applied AI Systems, Inc.
• BNR
• CAE
• CP
• CompEngServ
• DEC
• Hydro-Québec
• MDA
• NB Power
• Price Waterhouse
• Royal Bank
• Shell

Figure 1. A review of AI in Canada.

established as a nonprofit corporation, officially inaugurated in 1985. Membership in CRIM consists of 7 university members, 30 industrial members, and 57 employees. CRIM has designated artificial intelligence as a research priority in addition to research in other computer-based domains. CRIM has continued, therefore, to pursue fundamental AI research, while at the same time implementing a number of different applications. Research areas include continuing work in knowledge engineering, knowledge acquisition, testing and validation of expert systems, and the development of automatic speech recognition. Applications have included a family of expert systems to aid the hydroelectric industry with the diagnoses of electrical equipment for the purpose of maintenance and repair, collaboration in a robot project for the maintenance of high voltage transmission lines, expert knowledge on problems related to pitch deposits in kraft mills, plus others in meteorology, speech recognition, flight simulators, and business planning.

Information Technology Research Centre

The Information Technology Research Centre (ITRC) is one of several centers of excellence supported by the Ontario Premier's Council Technology Fund. The ITRC promotes fundamental and applied research in information technologies at the Universities of Toronto, Waterloo, Queen's, and Western Ontario. Research supported by the ITRC at these four universities includes the work of Professor Mylopoulos at Toronto on a system to facilitate the creation, access, and management of large databases, Professor Wong at Waterloo in knowledge acquisition, and Professor Wonham at Toronto and Professor Pang at Waterloo on the control of discrete-event systems. Biological and computational vision studies by Professor Tsotsos and others at Toronto include studies on the role of selective attention and parallel processing as evidenced in biological systems. Efficient neural net learning procedures, including the use of unsupervised learning procedures, are being developed in a project led by Professor Hinton, Toronto, with application to speech recognition and shape recognition. At Queen's University, Professor Browse and others are conducting investigations in both visual and tactile perception, while at The University of Western Ontario, Professor Pylyshyn and others are pursuing theoretical studies in cognitive science in areas where humans are extremely facile, such as tracking moving objects or perceiving the shape of 3-D objects, and also areas in which humans are very limited, such as performing mental arithmetic or detecting certain kinds of image changes. Work is also proceeding on image analyses applied to storm tracking on weather radar and in motion understanding.

PRECARN

Subsequently, and in a supplementary way, PRECARN was established as an industrial research consortium to conduct and promote long-term, precompetitive research and experimental development in advanced robotics and AI development. The name PRECARN is an acronym for PREcompetitive Applied Research Network, which indicates its purpose. It was a realization of the need to strengthen the "receptor capacity" in Canadian industries which led to the establishment of PRECARN as an activity complimentary to that of the CIAR. PRECARN's current membership consists of 34 Canadian corporations drawn from a broad cross-section of economic activities across Canada. Nine of their proposed projects have been selected for feasibility studies, with funding of $884 000 in total obtained from a combination of industry and government support. These include projects in telerobotics, mobile robots, the application of expert systems in advanced process control, and others. In addition, PRECARN has been instrumental in the establishment of the Institute for Robotics and Intelligent Systems (IRIS), a nation-wide research program in intelligent systems, managed by PRECARN and involving over 120 academic researchers, including Canada's top talent in AI and robotics. IRIS is one of 14 proposals selected for funding under the Federal Networks of Centres of Excellence Program.

PRECARN's IRIS Network

The Institute for Robotics and Intelligent Systems (IRIS) is one of 15 Research Networks recently selected by the Canadian Federal Government under the new $240 million Networks of Centres of Excellence Program.

IRIS is a collaborative effort between the academic research community and the private sector. It brings together within one nation-wide network the leading academic talents in the fields of artificial intelligence and advanced robotics, including the fellows of the Canadian Institute for Advanced Research. IRIS is managed by PRECARN, an industrial consortium for precompetitive research in AI and robotics, and is operated as a separate component to PRECARN's own industrial research network, but with the same procedures for project review and knowledge dissemination.

The research program of IRIS includes 22 projects within three areas of investigation: computational perception, knowledge-based systems, and intelligent robotic devices. The program will be carried out over four years, with a federal grant of $23 million. Over 100 researchers from across Canada will participate.

A description of the overall research program and its 22 projects is reprinted from the original IRIS submission.

Government Support Programs for AI R&D in Canada

In addition to the Natural Sciences and Engineering Research Council (NSERC), government support for AI technology development in Canadian industry is also provided through the industrial support programs of Industry, Science and Technology Canada (ISTC) and the National Research Council (NRC). Specific programs include the AI R&D fund, the Strategic Technologies Program (STP), the Microelectronics and Systems Development Program (MSDP), the Technology Outreach Program (TOP), the Defence Industry Productivity Program (DIPP) for which ISTC has responsibility, regional programs, and also the Industrial Research Assistance Program (IRAP) of NRC.

Natural Sciences and Engineering Research Council

The Natural Sciences and Engineering Research Council (NSERC) is a federal granting agency with the mandate to promote and assist research in the natural sciences and engineering. Total grants and scholarships will exceed $400 million in 1990–91. Some 17 000 individuals benefit from NSERC's research and training programs.

Support is provided for a wide variety of disciplines, including artificial intelligence, through a number of NSERC programs including the Operating Grants Program, the Strategic Grants Program, Research Partnerships, Scholarships, and International Programs. NSERC invested in excess of $3.6 million for artificial intelligence research in 1989–90.

AI in Universities, Government Departments, and Agencies

Canadian Space Agency

The Canadian Space Agency (CSA) supports a number of AI activities in the space sciences and applications, principally by the Space Station Mobile Servicing System (MSS) and Strategic Technologies for Enhanced Automation and Robotics (STEAR) programs.

The Canadian contribution to the International Space Station "Freedom" is the Mobile Servicing System (MSS). This is a collection of robotic elements configured to support the assembly, maintenance, and servicing of the space station. Major components of the MSS are the Space Station Remote Manipulator System (SSRMS), a large robotic manipulator arm with seven degrees of freedom; the Special Purpose Dextrous Manipulator (SPDM), a smaller two-arm robotic device; and the Mobile Servicing Centre (MSC), a mobile base for the MSS used to transport payloads around the station. The baseline capability of these components at the initial operating configuration of the station will incorporate strategic technologies, including expert systems utilizing AI techniques. In addition, these baseline capabilities have been designed with the objective of evolving toward autonomous operations.

Currently developed baseline capabilities include a command and programming language, a collision avoidance system, a target-based space vision system, and a procedural reasoning system for the capture of failure diagnosis knowledge that will execute in real time. A planning system is being developed to assist in the generation of plans for MSS operations. A fundamental component of the MSS Control Equipment is an Artificial Vision Unit (AVU), for both manual and automatic control of the manipulators and cameras. The AVU will support rendezvous and proximity operations, including payload tracking, payload capture, and payload berthing.

It is recognized that there will be a shift in emphasis from teleoperated robotics to autonomous robotics during the 30-year life-span of the Space Station Freedom.

The evolutionary capability of the MSS will require an increasing utilization of advanced AI techniques for which phase 1 preliminary feasibility studies are in place and phase 2 proof of concept projects are planned. These involve diagnosis, planning, tutoring, machine learning, vision systems, trajectory planning, and obstacle avoidance. Descriptions are given on the status of these, which, in most instances, involve projects contracted to the private sector.

Department of Communications

Canadian Workplace Automation Research Centre

The Centre, forming part of the research program of the Department of Communications, was opened in 1985 and maintains a focus on office automation. The CWARC acts as a catalyst for the growing use of telecommunications and computer technologies in response to the needs of knowledge workers in the office of the future. Within CWARC, most AI projects are undertaken in the Advanced Technology Directorate which consists of two groups: the Computer Aided Translation Group and the Expert Systems Group.

A short-term objective of the former is the development of a translator's workstation, of which a modest first version based entirely on commercially available software is currently in trial use. In addition, a project is about to be launched in bilingual text generation which will automate the drafting of several types of bilingual summary reports for Statistics Canada in both target languages. A third component of the machine translation program has resulted in a prototype system for the translation of agricultural market reports. It is the first machine translation system to be completely reversible, that is, it can translate from French to English, or English to French, using the same grammars and dictionaries.

Work in expert systems includes systems for knowledge management (knowledge acquisition and knowledge modeling), systems for education and training, and information management systems, such as those for automatic indexation of documents. Expert systems completed to date include one to aid in the evaluation of project proposals submitted for funding to the Economic Regional Development Assistance (ERDA) Program. Another, called AMETHYST, assists pay and benefits clerks and consists of 13 expert systems, each corresponding to an area of pay and benefits (e.g., salary eligibility, medical insurance, income tax, etc.).

Communications Research Centre

Research projects in two areas at the Communications Research Center fall within the domain of artificial intelligence. These are the work on speech recognition using neural networks and systems enabling language conversations with computers.

Research on automatic speech recognition is aimed towards speaker-independent recognition of continuous speech where many potential applications reside. The Conversational Hypertext Access Technology (CHAT) allows people to retrieve paragraphs of information about a specific topic in response to questions asked in either English or French.

Department of National Defence

The Department of National Defence (DND) regards AI as a technological area to be investigated as a methodology capable of assisting in particular classes of problems. As AI techniques become more widespread and mature it is expected that AI technology will find its way into defence research applications and eventually into operational situations. Meanwhile, National Defence Establishments in Dartmouth, Valcartier, Ottawa, Downsview, Suffield, and Esquimalt will continue to monitor the technology and develop specific applications where appropriate. Current work involves the application of AI and related techniques in vessel detection and classification, battle and situation knowledge, information fusion, threat evaluation and weapon assignment, radar and communication systems, computer vision, model-based object recognition, robot-based ordinance detection and reconnaissance, image processing, signal detection, budget planning, neural networks, and aids to emergency response managers.

Energy, Mines and Resources Canada

In the execution of its mandate, Energy, Mines and Resources Canada (EMR) conducts research involving artificial intelligence in a wide variety of application areas. These range from photo interpretation and satellite data processing at the Canada Centre for Remote Sensing to the production of maps at the Canada Centre for Mapping, advanced process control and other applications at the Canada Centre for Mineral and Energy Technology (CANMET) and discovery methodologies for geochemical and mineral deposits, including gold, at the Geological Survey of Canada.

AI activities at the Canada Centre for Remote Sensing (CCRS) are very extensive. Artificial Intelligence is seen as an essential technology for the processing of remote sensing and to solve the increasingly complex data fusion problems. Expertise includes the design and use of RESHELL and creation of the first prototype for a hierarchical expert system to aid in the analysis of LANDSAT data. A second hierarchical expert system, called System of Hierarchical Experts for Resource Inventories (SHERI) is now underway. SHERI will contain more than 150 expert systems and is scheduled for completion in 1993.

The production of maps involves the transfer of data from layer scale to smaller scale maps. Previous efforts at the Canada Centre for Mapping (CCM) concentrated mainly in the graphical manipulation of object representation. Research objectives now include developing the rule base to provide a comprehensive approach to generalization in the digital environment.

Other AI applications include a Statistical Mapping Advisor (SMA) and the investigation of automated name placement for topographical map production.

In CANMET, the Mineral Science Laboratories (MSL) have completed three projects on expert systems applications in mineral processing. Offline process control applications have included a clinker grinding expert system developed in conjunction with Cement Lafarge and a second system for operator assistance in the control of a zinc leaching and refining plant developed in collaboration with a team from Lavalin and Canadian Electrolytic Zinc (CEZ). Experience with expert systems for on-line process control, which are few in number, has been gained in a joint project with Kidd Creek Mines in a system for X-ray analyzer data validation, followed by data substitution when the essay data are deemed to be erroneous.

Experience gained reveals a need for expert systems to be closely integrated with existing computer hardware and software, and a need for integration of, or conversion between, symbolic and numerical processing. These needs and trends, while absorbing new technologies such as expert systems into the main stream of industrially based activities, at the same time tend to reduce their external visibility.

Other CANMET applications include rock stress analysis and knowledge-based ground support systems at the Mining Research Laboratories and the prediction of residual fatigue life for welded structures at the Metals Technology Laboratories in addition to expert systems for the selection of metal alloys, the control of elective furnaces, mould cavity design, and foundry operations. The Coal Research Laboratories are developing a prototype expert system related to the recovery of coal fines, the Energy Research Laboratories are developing an expert system for assessing the suitability of any coal for use in coal gasification, while the Efficiency and Alternative Energy Technology Branch is developing three expert systems related to building energy management systems and the design of energy efficient buildings.

The Geological Survey of Canada has a number of projects involving production rule expert systems and one under consideration involving neural nets. Several are related to Geographical Information Systems (GIS) in order to undertake mineral resource assessment tasks. Projects involve the identification of gold deposits, the forecasting of geomagnetic storms, the classification of metamorphic rocks, and the use of hypertext in generating a users' manual for oceanographic survey equipment.

Government Consulting Group

The Government Consulting Group (GCG), a successor organization to the Bureau of Management Consulting (BMC), continues the previous commitment and active interest in Strategic Technologies, including AI as a major component. Projects in AI and related disciplines have included those for automatic speech recognition (ASR), data fusion involving neural network selection, the recognition of sonar signals by neural networks, naval warfare gaming and simulation, expert systems for toxicity assays, plus state-of-the-art surveys and educational seminars.

National Research Council

At the National Research Council (NRC) the Laboratory for Intelligent Systems in the Division of Electrical Engineering began its program in applied Artificial Intelligence in early 1987. The activities of the program are directed towards the development of knowledge-based systems tools and methodologies and their application in response to specific requirements. The general objective of the Laboratory is to provide, through applied research and collaboration with Canadian industry and universities, intelligent systems technology which will make Canadian industry more competitive in the emerging fields of sensor-based robotics and knowledge-based systems.

Some 60% of NRC resources devoted to AI are located in the Information Technology Section where work is focused on the development of generic AI technologies and methodologies with an application orientation and in collaboration with others. Projects include a natural language interface which functions in the French language, a machine learning system oriented to manufacturing process planning, a jet engine maintenance and diagnostic system, a realtime expert system for process control applications, and an intelligent advisor system to provide assistance to database users.

AI projects and activities in other areas of the NRC include work on robotics path planning, autonomous vehicle development, procedures for resistance welding, medical diagnosis, electromagnetic hardening (interference), a maintenance advisor for a paper making machine control system, the use of industrial materials, and developments associated with Canada's participation in the Space Program. In addition, industrial projects in AI are supported through NRC's Industrial Research Assistance Program (IRAP), and a special group dealing with information technology in the construction industry has been established in Calgary as part of NRC's Institute for Research in Construction in collaboration with the Alberta Research Council.

Provincial Research Organizations

A review of AI activities among the members of the Association of Provincial Research Organizations (APRO) indicates AI activity within a number of locations. At the New Brunswick Research and Productivity Council (RPC) researchers are using fuzzy logic to implement some portions of a control/alarm/advisory system for draglines used in open-pit mining. At the Montreal laboratory of the Centre de recherche industrielle du Québec (CRIQ) current activities are in machine vision, automated inspection, and robotics. Image analysis techniques have also been used to identify parasites in fillets of fresh fish and recently work has begun using a neural net simulator in a character recognition application.

AI activities appear to be particularly strong at the Alberta Research Council where the Advanced Technologies Department was established in 1984, now renamed as the Advanced Computing and Engineering Department and having more than 40 employees. Since established, the department has conducted over 40 projects in AI, robotics, and automation with contract revenue in 1989 of $1.8 million. The department is organized into five technology groups:

- Artificial Intelligence (AI) in Operations Group
- Learning and Collaboration Group
- Pattern Analysis and Machine Intelligence Group
- Realtime Artificial Intelligence Group
- Robotics and Automation Technologies Group

Projects undertaken for industrial companies and others have covered a wide range of application areas including telephone switching, highway design, automated scheduling of a fermentation plant, a learning system that included curriculum design and the generation of examinations, an expert system for well log analysis, and a character/word recognition system to interpret handwritten digits and symbols on maps and map overlays, maintenance and fault diagnosis in the CF-18 aircraft, ballast control for semi-submersible oil drilling units, the management of spills of hazardous materials, severe weather forecasting, and others.

Transport Canada

Since the transportation sector is technologically intensive, the operational responsibilities of Transport Canada require assessment of the feasibility of applying new technologies and supporting their application where appropriate. Examples include air systems navigation and regulatory responsibilities for safety and security. In total, the department contracts out over 95% of its R&D program which amounts to $25–35 million annually.

Starting from its initial AI research and development activities in 1983, the development and implementation of relatively complex expert systems is now widespread throughout the department, and becoming more so in the transportation sector as a whole. Early assessment indicated that virtually every transportation operation and regulatory task can be effectively and economically supported by an expert system. Potential areas for application were established such as:

- Vehicle Crews
- Vehicle Maintenance
- External Control of Vehicles in Transit
- Permanent Way Control and Operations
- Terminal Control and Operations
- Operational Planning and Regulation
- Strategic Planning

each of which can be further subdivided; for example, Vehicle Crews includes on-board navigation and the provision of emergency advice while Permanent Way Control and Operations includes weather forecasting. By 1988, when some 15 expert systems had been developed or were approaching completion, it had become evident that the results of the department's AI R&D program were exceeding all expectations and a new series of application and investigative projects were planned for which details are provided in the references cited.

Currently Transport Canada is spending over $2 million annually on AI R&D, primarily for expert system development. In addition the Aviation Group has established an AI testing and evaluation facility for the department. Exploratory R&D in neural nets, voice recognition, and intelligent tutors suggests that these technologies are approaching the point where they can be applied in transportation.

Universities

A general review of AI activities in academic institutions in Canada traces development from its early beginnings through the growth patterns of the last decade to its present capabilities. Key events include the formation of the Canadian Society for Computational Studies of Intelligence/Société canadienne pour l'étude de l'intelligence par ordinateur (CSCSIO/SCEIO), the biennial conferences, the establishment of *Canadian Artificial Intelligence* as a quarterly publication and *Computational Intelligence* as a research journal, plus the CIAR project and the even more recent IRIS project. Case descriptions, at a more detailed level, are given for a few selected universities.

University of Calgary

The Knowledge Science Institute was established at the University of Calgary in 1985 with the objective of performing studies in knowledge science, technology, economy, and social impact and to promote interdisciplinary research on knowledge-based systems. A strong network of activities has been created, involving researchers from the computer science faculty, other university departments, visiting experts from other countries, and industry. Research in knowledge engineering has progressed in many key areas, including a study of past and future trends in the computing industry and the automation of knowledge acquisition techniques for the rapid prototyping of knowledge-based systems. One area of application is in the area of sports where the acquisition of playing skills is of interest to both individual players and coaches.

A series of workshops on knowledge acquisition, begun in 1986, has resulted in great interest worldwide and has made a significant impact to the literature on this subject. Other studies include research in machine learning and programming by example, resulting in a robot-based drug-dispensing system for a hospital and research in system theory and organization with a capability in text compression methods and methodologies for the management of uncertainty in knowledge-based systems. A major part of the research centers around logic programming, which includes work on the use of interval arithmetic for AI programming languages and on parallelism in logic programming.

University of Guelph

An outline of the activities of three groups working in AI-related areas within the Department of Computing and Information Science at the University of Guelph is given. Currently there is much research attention on the use of parallel and pipeline architectures for image processing applications. Pipelined systems are increasingly being used in industrial vision systems, particularly in manufacturing visual inspection. Work at the University of Guelph focuses on the software programming support necessary for such systems, particularly when open access to the software is required, as in a research environment. Work in pattern analysis and machine intelligence is being applied in two areas, namely the analysis and synthesis of patterns from molecular sequences, such as RNA and tRNA sequences, and a pattern recognition approach to the analysis of historical earthquake data in China. Work on theoretical issues in uncertainty issues investigates ideas in probabilistic logic, nonmonotonic reasoning, Dempster-Shafer theory, and the ability of probabilistic logic to deal with either small or large probability values. Applications in medical diagnostic systems are being developed in veterinary medicine as in equine colic, a complaint of horses often requiring emergency surgery, and liver disease in small animals, plus a natural language system developed to handle the learned vocabulary and grammar of dolphins.

Queen's University

Artificial Intelligence is addressed at Queen's University in an active program of both tracking and research which is administered jointly by the Departments of Psychology and Computing and Information Science. Research in AI is concentrated in the areas of robotics, perception, knowledge-based systems, and music.

In robotics and perception, work is directed towards the problem of image segmentation as required for robot vision systems. A number of techniques are being explored, including discrimination based on object texture since it is known that humans can manually explore and identify objects based on shape, hardness, compliance, texture, and other tactile sensory data. Work in knowledge-based systems involves the Nested Interactive Array Language (Nial). A principal application of Nial is in areas where the powerful operations for dynamic manipulation of both symbolic and numeric data permit straightforward implementation of AI techniques. Development of the Nial AI toolkit is on-going, with applications of the programming language, at Queen's for the space station Mobile Servicing System plus a project for the Bell Canada capital recovery group. Nial has also been used in external applications for aircraft maintenance, well logging, sewage system control, software testing, and insurance policy screening. Research work on AI and music involves symbolic logic as applied to knowledge-based modeling. This includes a "composer's assistant" as a tool for music composition and a programming language for composers. A study of methods for extracting the semantics of two-dimensional data representations, such as graphs, charts, or wiring diagrams, uses sheet music as an initial example for the development of general representation and interpretation techniques.

University of Saskatchewan

AI is a vigorous and important part of the Department of Computational Science where 5 of the 16 full-time faculty members have their primary research interests in this field. A wide variety of topics are covered, including representation of knowledge, natural language understanding, dynamic planning, cognitive modeling, qualitative reasoning, probabilistic reasoning, rule-based expert systems, shape analysis, segmentation, motion analysis, human factors, and others. The Laboratory for Advanced Research in Intelligent Educational Systems (ARIES) was founded in 1987 to provide a focus for projects concerned with applying AI techniques to training and education. The Student Computing Environment (SCENT) project aims to create a supportive and intelligent programming environment for use by novice programmers, including recognition of student coding strategies and misconceived strategies.

Projects in image processing include the development of techniques that can be used in the recognition of shapes that are deformed under controlled conditions and the understanding of human body motion, while those concerned with uncertainty management in expert systems include the use of Bayes' theorem, Dempster-Shafer theory, and fuzzy logic as useful approaches.

Simon Fraser University

At Simon Fraser University (SFU) an emphasis on interdisciplinary studies has provided a positive structure and support for AI research in the five areas of expert systems, computational vision, computational linguistics, knowledge representation, and learning systems. Application areas include realtime process control, natural language understanding, automatic translation, studies in color constancy and vision, robot motion planning, intelligent extensions to CAD systems, and the integration of multisensory data for robot vision systems.

University of Toronto

Research in biological and computational vision at the University of Toronto is an interdisciplinary activity spanning several departments and affiliated with activities at York University. Following early Ph.D. thesis dissertations in this area, a formal vision research group in the Department of Computer Science was established in 1980, with collaboration and cooperation in other departments following. The research is guided by the philosophy that to achieve deep and meaningful results one must take seriously the hypotheses that human visual information processing can indeed be modeled computationally. This requires, among other considerations, that all biological observations should be carefully considered, that all appropriate mathematical tools must be utilized, and that theories produced must also have a predictive impact on the experimentation of biological scientists in a give and take interplay.

Projects, as of January 1988, conducted by a faculty–student research team of substantial size, include, among others:

- Early motion sensing and measurement
- Intermediate representations for images
- Models of image structure and primitive features
- Early chromatic processing
- Visual stereoscopic computation
- Extracting motion and structure from images
- Self-organizing feature maps
- Root vision and navigation

Application Development and Use in Business and Industry

Acquired Intelligence Inc

Acquired Intelligence Inc is a privately held, Canadian-owned corporation involved in the research, development, and commercialization of artificial intelligence. Although small, it is perhaps representative of other small- and medium-sized enterprises in the high technology sector of private industry. Its relatively small staff include five Ph.D.'s, four M.Sc.'s, and several B.Sc.'s with backgrounds in computer science, cognitive psychology, engineering, neuropsychology, economics, and geology. The company's products and services are strongly oriented to AI software tools related to knowledge acquisition. These run on readily available personal computers and enable users to develop and maintain expert systems with a minimum of specialized human assistance. Use and testing of the system has included applications in haematology, forestry, mining, market analysis, medicine, space, agriculture, tutoring, neuropsychology, and the construction of knowledge bases from written materials such as city parking regulations.

Applied AI Systems, Inc

Applied AI Systems, Inc provides expertise in the research, design, implementation, and deployment of practical AI systems, education, training, and business consultation. Studies and reports listed include planning and state-of-the-art surveys evaluating AI technologies and opportunities for the transportation sector, voice interactive technology, the automation of marine operations, reviews of AI research in Japan, Britain, and Europe, natural language understanding in neural networks, vehicle dispatching, and others.

Bell-Northern Research

As the research and development organization for Northern Telecom, BNR has a staff of more than 6700 people in laboratories across Canada, the United States, and the United Kingdom. The scope and focus of the AI program, which has been running for several years, is to explore and develop AI techniques for use in products and services provided by BNR and its parent companies, and for use in design and development processes. The complexity of telecommunications networks is increasing rapidly with the introduction of new functions and services, increased connectivity, and the advent of new components. At BNR it is believed that many of the tasks involved in establishing and operating such networks can be done using expert systems and, in fact, that such systems will soon be essential. Applications of AI at BNR include systems for digital telephone switch maintenance, advice to customers on loop access, transforming business user specifications into detailed hardware, software, and layout descriptions, circuit design engineering, software engineering, and natural language interfaces for database queries.

CAE Electronics Ltd.

Activities in artificial intelligence within CAE Electronics Ltd. began in 1985 with an evaluation of AI's potential contribution to the company's needs and interests. In June 1988 a specialized group was formed to conduct existing projects and pursue new applications. Projects include participation, with others, in the PRECARN project to develop an Advanced Process Analysis and Control System (APACS). Briefly, the project, as an extension to the current state-of-the-art in computer-based process control systems, is oriented to advances in two technical areas, namely the use of knowledge-based systems to assist nuclear power plant operators in technical decision-making and the possibility of developing "faster than realtime" simulation capabilities supporting on-line, realtime diagnosis of systems faults and the prediction of plant behavior.

The FLITE project is a program of research in the application of AI in the domain of aerodynamic simulation. The problem arises in the tuning of the aerodynamic models used to represent aircraft in simulators manufactured by CAE. Typically a hundred or more distinct flight-test manoeuvres are tested and recorded, to which the output of the aerodynamic model used in the simulator must be closely matched by adjustment of model coefficients and simulator inputs. The user interface, mathematical optimization, qualitative reasoning, and automatic parameter tuning capabilities of the FLITE system provide considerable assistance in this process. Also in the simulation domain, a system known as ITEMS provides user assistance to military customers in the use of tactical mission simulation systems for training purposes. Knowledge-based techniques are also being used to increase the recognition accuracy of speaker-dependent, continuous speech recognition systems oriented currently towards air traffic control applications.

Canadian Pacific

Applications implemented at Canadian Pacific, since AI technology use began in 1984, include an expert system used to troubleshoot 80% of the faults occurring in mobile radios, a maintenance advisor based on the spectrographic or physical analyses of locomotive oil samples, and a computer-assisted train dispatching system. The latter system, now widely used throughout North America, applies train dispatcher rules and lists all valid choices to support dispatching. This reduces the work necessary, reduces dispatcher stress, and increases rail safety. Other applications at CP include an advisor for determining when rail cars should be accepted from other railways, and when they should be returned according to the regulations of the American Association of Railroads, and a system that permits entry of database queries in English. New applications under development will consolidate the knowledge of experienced mechanics in an expert system for diesel engine maintenance and repair, a system for scheduling track maintenance while minimizing train delays, and an intelligent reporting system that will extract information from company reports, thereby providing the user with needed information on an exception basis only.

CompEngServ Ltd.

CompEngServ, as a 15-year-old company dedicated to advanced computer hardware and software, has been actively involved since 1983 in developing AI solutions to problems. These have included the classical AI disciplines of expert systems, rule-based systems, planning systems, vision systems, robotics, and neural networks as well as AI tools such as object-oriented programming languages.

The short summary provided of major projects includes an expert system for air traffic conflict resolution (CORES), implemented using the LISP-based system builder KEE, and EMIR, an electromagnetic interference expert system. A marine environment data management system (EDMAN) is an intelligent interface to archival data developed for the Department of Fisheries and Oceans. Use can include ensuring the quality of data before they are entered into the database as well as providing assistance to users of the database.

Digital Equipment Corporation

Digital Equipment Corporation, as an early user of expert systems technology, maintains an aggressive program of expert system development and a staff of 600 AI specialists, world wide. Savings from deployed expert systems are estimated at $200 million annually and growing at 40% per year. Separate groups in the AI Technology Center (AITC) are organized to focus on AI product development, internal applications, external consulting, and AI research.

Recognizing the need within the company for a large number of skilled knowledge engineers, more than 450 "students" have graduated from the company's Knowledge Engineering Curriculum, later made available to selected customers as the AI Fellowship Program. Employees of more than 80 major customer accounts have graduated from this program, including 14 from Canada. Cooperative expert system projects with Canadian customers are listed.

Recent trends to reduce the cost of expert systems include the identification of solutions that are common to all or most companies in an industry group and to classify expert systems according to the nature of the problem solved (diagnostics, configuration, scheduling, etc.) rather than classifying systems by industry.

Hydro-Québec

Because power system activities are knowledge intensive and must achieve a high degree of reliability and efficiency, they must use the most powerful and promising technologies available. This includes artificial intelligence, and expert systems in particular. At Hydro-Québec technology management methodologies have been developed, from an applications point of view, to determine what expert systems are possible, to eliminate undesirable candidates, and to rank the promising applications according to their economic and other benefits. These benefits include the need to transfer expertise to a large group of users, to protect expertise that is disappearing, to facilitate a need for training and a need for expertise in hazardous environments. Applications described include expert systems:

- for preliminary design of transmission systems
- for interpreting events in transmission systems
- to improve the processing of alarms and status changes
- for diagnosis of large power transformers
- for diagnosing the safety of earth dams
- used to simulate the operation of hydroelectric complexes
- linked to a CAD system, to produce layouts for 25 kV substations.
- for the management of energy exchanges with neighboring power systems to maximize sales revenues and minimize costs

MacDonald Dettwiler and Associates Ltd.

MacDonald Dettwiler and Associates Ltd. is a leader in the technologies of remote sensing, particularly in the reception and processing of imagery from space. This includes systems for processing data from earth resource satellites such as LANDSAT and SPOT, from meteorological satellites, and from synthetic aperture radar (SAR) satellites.

Application examples of expert systems cited include:

- a meteorologist's assistant for the 12-hour prediction of precipitation achieved from the automatic projection of the future positions of warm and cold fronts on maps

- a system for making same-day predictions of severe weather in the Canadian prairie provinces

- a knowledge-based system for automatic extraction of planimetric features from remote sensing imagery based on knowledge associated with roads

- a knowledge-based system, developed in conjunction with MacMillan Bloedel Ltd., for predicting faults and planning preventative maintenance for large off-highway logging trucks. The system uses either recorded data from specially installed sensors or observations collected by a human user.

New Brunswick Power Commission

New Brunswick Power has undertaken a number of expert system applications directed towards the solution of specific problems encountered in the utility. A realtime knowledge-based controller has been developed in Prolog in collaboration with the University of New Brunswick with the objective of minimizing the effect of failure-induced response in electrohydraulic governors on hydro generating units. A second project examined the use of rule-based alternatives for the front end of a software system used for unit commitment. This is part of an operational hydrothermal scheduling and river modeling software package. Other applications have involved the use of expert systems in distribution systems where fault diagnosis and power restoration are an area of major concern, a prototype expert system for automatically generating test data for validating the moderator temperature control program in a nuclear power plant, a knowledge-based system to assist in the selection of channels when refuelling the reactor in the Point Lepreau Nuclear Generating Station, and a computer program to aid in the analysis of established emergency operating procedures.

Price Waterhouse

As perhaps typical of other large consulting firms, Price Waterhouse has identified expert systems as one of the core technologies within its Information Technology Practise. The special Canadian unit for emerging technologies maintains close contact with similar groups within the company in the United States, the United Kingdom, and Australia.

The list of knowledge-based and expert-system-based projects reflects the diversity of AI applications and their growing penetration in business, industry, and government. These include expert systems to screen proposals for the creation of new standards, risk assessment of medical devices, assessment of the impact of new technologies on the containment of healthcare costs, an advisory system for chemotherapy treatment of cancer, and educational seminars on AI for potential users. Trends and future directions foreseen include broadening and integrating the application of the technology into auditing and tax practise, systems management, and conventional information systems.

Royal Bank of Canada

Led by projects carried out in its Technology Planning and Development Department, the Royal Bank of Canada is continuing to adopt AI technology, mostly on the use of expert systems to solve particular types of problems. Of particular note is the use of expert systems which make it possible to capture the business knowledge of experienced staff, and then make this available to others for use in traditional areas such as expert advisers, operational help desks, and risk assessment systems.

Shell Canada

Over a period of four years expert system technology has advanced from being largely unproven in the oil industry to a position today where it is being used at Shell Canada as an accepted technology to solve a class of problems that are very difficult, or even impossible, to solve with conventional computer-based techniques. The expectation or vision for expert systems has included problems that are normally solved by heuristics, judgement, and rules of thumb. A second class of applications is in commercial areas such as tax rules, royalty rules, and personnel policy. Conventional computer systems in these areas are expensive to maintain and subject to unreliability because of frequent changes. Expert systems technology offers a better way to build these as rule-based systems. A third application area is directed towards capturing the knowledge and expertise of senior people in technical field operations, financial operations, personnel and pension policy, and passing this experience on to less experienced staff. To date approximately a dozen systems have been built in all three areas, generally using commercially available expert system shells. These provide advice to drilling engineers, process operators, toxicologists, credit risk evaluators, and realtime support for a help desk operator.

Conclusion

A conclusion that can be reached is that while artificial intelligence is still a young discipline, its development and application is continuing as an emerging field. It is also pervasive and highly interdisciplinary. While most of the work in Canada is research based and oriented to future applications, a healthy application base is emerging that would have been difficult to identify just a few years ago. With over 1000 persons reported as working full time in AI it can be said that AI is alive and well. The AI research community in Canada is highly respected in international circles. New and novel applications continue to be pursued vigorously in business and industry.

The NRC Associate Committee on Artificial Intelligence has chosen as its top priority the encouragement of closer collaboration between AI users and the AI research community in university and government laboratories. It is with this objective in mind that the current report has been prepared.

More complete descriptions are found throughout the remainder of the report, supported by a detailed cross index at the end. Once again, the committee wishes to express appreciation to the many organizations that have contributed to its preparation.

Section 2

Introduction

Introduction

Before one can begin to consider the business, technical, or social ramifications of artificial intelligence, it is necessary to have some reasonable definition of what is meant by the term "artificial intelligence."

Artificial intelligence (AI) is not easily defined. It may be considered to be founded on a branch of computer science that concerns the conception, design, development, and application of computer-based systems to perform tasks that require the exercise of intelligence as exemplified by capacity to recognize, reason, decide, and learn. Thus, application of AI seeks to incorporate such human-like skills and expertise into computer-based systems. These systems include what are called knowledge-based and expert systems.

It is an essential characteristic of an application of AI that the associated computer system has provision for representing definable elements of knowledge about the context in which the system will be used and about the goals intended to be realized in its use. Further, when the range of conditions of application of the system is appropriately limited, the behavior exhibited by the system is at least intended to be rationally attributable to the content within it of the representation of its knowledge and goals. These characteristics of AI systems will be illustrated by a variety of examples.

Computer systems which realize artificial intelligence commonly display advanced forms of input and output such as:

- Speech input and voice recognition
- Natural language processing
- Machine vision
- Photo interpretation
- Image understanding
- Advanced forms of sensor-based input
- Voice synthesis and speech output
- Mechanical and manipulation output capability (as in robotics)

Applications of AI include, but are not limited to, the following:

- Natural language processing

- Machine translation

- Expert systems, such as those used in the diagnosis and treatment of disease, monitoring and diagnosing faults in equipment, industrial plants, telecommunications, and electric utility systems

- Operations scheduling and process planning in manufacturing

- Robot and other autonomous systems with vision or other sensing capabilities that permit them to respond to changes in their external environment

Expert systems are of particular interest because in terms of application they represent a leading edge of AI. They usually exhibit most, but not necessarily all, of the attributes listed in Table 1.

A short list of potential application areas is shown in Table 2 in order to illustrate the breadth and pervasive nature of artificial intelligence. It is necessary to recognize this breadth in order to place the subject of AI in perspective.

Basically the benefits of AI tend to be derived in the following three ways:

1. Knowledge-based systems tend to capture, represent, and distribute knowledge more uniformly throughout organizations and societal structures. Industrial firms, for example, find that the knowledge and experience gained by long-term employees can be retained and distributed throughout the organization in a more uniform and accessible way than was previously possible.

2. Improved decision making is achieved via the quality of the knowledge base and reasoning process employed, and also in the speed of response. Systems are already being developed for a wide variety of applications, particularly in the form of expert systems.

3. As a result of improved and more natural interfaces and means of communication between people, between systems, and between people and computers. The possibility of natural language processing, the possibility of improved methods of machine translation for the many written and spoken languages of the world, the possibilities envisaged from computer recognition of speech and machine vision are all examples of this nature.

Table 1. Expert system attributes.

Nature or application and use	Expert systems are used to provide advice in much the same way that it would be obtained and used if a human expert were available.
Domain expert	In the building of an expert system it is necessary to have available, and to draw on, the knowledge and experience of a person who is already an expert in the subject area. This person is referred to as the domain expert.
Rule based	Rather than being based on procedures written in algorithmic form, many expert systems derive their knowledge from a collection of rules.
Method of knowledge representation	The system knowledge is contained and given by the rules in the knowledge base rather than being represented by procedures given by computer coding as in "conventional" programs.
Symbolic representation	Expert systems tend to work with information in symbolic form rather than being based on calculations performed in numeric form. Use may be made of input output or rules expressed in a natural language such as English, rather than being expressed numerically.
Search strategies	In seeking a problem solution, expert systems examine, select, and manipulate rules in the knowledge base by means of search strategies such as forward and backward chaining. The portion of the system performing the logic function in an expert system is sometimes known as the inference engine.
Computer language	Expert systems are often written in list processing languages such as LISP, or logic programming languages such as Prolog rather than in procedurally oriented languages such as Fortran.
Traceability	An important attribute of expert systems is their ability to explain how a given result was obtained. Also known as transparency or explanation facility.
Apprenticeship	There is often a need to subject expert systems to a period of apprenticeship as part of their development because it is harder to verify and validate expert systems than it is for computer programs of a numerical or computational nature.
Designed to grow incrementally	Since the knowledge base is separate from the search logic it should be possible to add new rules or modify existing rules without having to review or reorganize the original structure.

Table 2. Application areas for AI and expert systems (1).

Application areas	Expert systems	Application areas	Expert systems
Manufacturing	Design synthesis and optimization Process planning Process parameter selection Robotic vision and intelligent robot systems	Education	Intelligent tutoring AI-based authoring environments Discovery environments Educational interfaces
Process control	Analysis and reduction of alarm conditions Defining process management rules Modeling the performance of best operators	Career guidance	Advising on education paths needed for employment desired in specified occupations
Materials	Selection of best materials for a given design function Selection of coating materials and finishes	Transportation	Air traffic control Equipment maintenance Inventory control
Resource industries	Evaluation of exploration data Crop and herd management in agriculture Mineral prospecting Mine planning and design Operations management	Financial planning	Advice on investment decisions and planning
Maintenance	Diagnosis of faults in plant and office equipment	Space	Remote sensing and data interpretation
Environment	General weather forecasting Prediction of hail storms	Health care	Medical diagnosis Determining optimum treatment cycles, e.g., as in radiation for cancer
		Automated language translation	
		Speech recognition	
		Natural language Processing	
		Military battle management	

Although definitions exist for what is meant by "artificial intelligence," examples constitute one of the best ways by which to gain an understanding of it. The publication *Canadian Artificial Intelligence* (see Appendix 2) regularly publishes a *World Watch* containing brief summaries of papers published in the field of AI. These describe many of the concepts, current developments, and application areas of AI, and illustrate its widespread and pervasive nature. The examples published in 1989–90 with emphasis on applications, explain much of what AI people plan, propose, and do.

At the same time these illustrate the many sources from which the benefits of AI will be derived, by making the knowledge and experience of experts more widely available, as in expert systems, and by making computers more accessible, as in the many forms of new input/output interface.

References

1. Scrimgeour, J. Expert Systems — A Review of their Potential and Application in Canada. *Canadian Conference on Industrial Computer Systems.* May 1986. (Principal source.)

Section 3

Consortia, Collaborative and Funding Organizations

Canadian Institute for Advanced Research

Keywords: Artificial intelligence, autonomous agents, Canadian Institute for Advanced Research, cognitive science, computational architecture, computational learning, computational linguistics, computational logic, computational vision, control theory, cross-institutional interaction, fellowships, graduate student meetings, interdisciplinary training, interdisciplinary workshops, knowledge representation, neuroscience, scholarships, telerobotics

The Field of Artificial Intelligence

Artificial intelligence has been defined as the study and design of systems which can exhibit behavior that would be considered intelligent if performed by a human. This, of course, is much too general. It not only leans on our intuitive understanding of "intelligence," but also includes all sorts of behavior that might later turn out to have been done by "unintelligent" means, but which nonetheless might impress people (such as judging the mass of a distant star, which might be done by a complex inference, but might also be done by some highly sensitive new instrument which does no intelligent processing at all). One of the tasks of the early stages of a science is the discovery of its natural boundaries. Many people now believe that the mark of intelligence is the ability to manipulate knowledge; that which makes some systems (like humans) capable of exhibiting intelligence is their ability to acquire, store, manipulate, and draw inferences from a large corpus of knowledge, and to decide upon a course of action based on this knowledge as well as certain goals or utilities. This general conception of intelligence presents a framework for understanding some of the apparently disparate streams of research that are carried out under the one heading of artificial intelligence (AI).

There are three main streams of research focus in AI, paralleling some of the aspects of intelligence mentioned above. The first stream is concerned with the acquisition of knowledge from the world. This involves the study of perception, by which knowledge is generally acquired from our natural environment, and of learning — an area that has made rather poor progress within psychology in the last half-century but which has recently seen a revival of interest among AI researchers.

The second stream concerns the study of ways of representing and manipulating knowledge, for purposes of reasoning and drawing new inferences from it. This stream is closely allied with the study of symbolic logic and has developed a number of formal calculi for representing various kinds of knowledge. This includes logics for representing and reasoning about beliefs, goals, conditionals, time, modalities (such as "should," "must") and other kinds of conceptual contents. Knowledge is manipulated by certain kinds of rules embodied in computer programs. In some way these rules resemble rules of logic, although they are typically quite different from the usual deductive logics with which most people are familiar (e.g., in syllogistic reasoning). For example, they may not be strictly truth-preserving but may use plausible heuristics, they may be probabilistics, they may make assumptions that are later withdrawn as new evidence is encountered, and so on.

The third stream is concerned with converting decisions into real actions. It studies the problem of planning a sequence of executable actions for achieving some goal, of managing perceptual-motor coordination, and of actually moving physical limbs or other actuators to carry out actions.

Throughout these various pursuits, there are always two kinds of problems that must be addressed. The first of these is "symbolic." It is concerned with problems of encoding the percepts or knowledge in terms of some internally stored codes and of manipulating these symbolic codes in terms of certain rules. It is not only the clear cases of reasoning and problem solving that involve such processes, but there are also good reasons for believing that they must go on in vision, in memory retrieval, in understanding sentences, and in other tasks where we have no awareness of deliberate reasoning. As well as attempting to discover encoding systems and rules that can carry out the task (e.g., to understand a sentence in English or to understand a LANDSAT image), AI researchers are also very interested in certain complexity measures associated with the process, and worry about ways to prevent the dreaded "combinatorial explosion" wherein the number of steps required to carry out the task may increase exponentially with the size of the input. The study of the symbolic aspects of AI often fall within the discipline of computer science, but they may also involve mathematical logic, linguistics, and the study of problem-solving and decision-making methods used by humans — which are typically the concern of psychology, economics, or even anthropology.

The second kind of problem arises because many AI researchers want eventually to design real physical systems which can carry out the processes under study. Because of this, an important aspect of AI concerns designing mechanisms (arrangements of components) that can carry out the relevant processes. This does not just mean the design of electronic hardware; it also involves the more abstract study of systems that implement intelligent processes. For example, it involves the study of what are called computational architectures, or classes of computational structures, such as highly parallel machine designs. It also involves the design of all physically realizable mechanisms that are needed to enable the intelligent activity to occur, including sensors, effectors, different types of memory stores, and control structures. In the case of the study of autonomous agents ("robots") or semiautonomous systems ("telerobots") a great deal of research has to go into the instrumentation side of technology; new materials may have to be exploited (especially for designing effectors and manipulators) and new systems of sensors may have to be developed.

The design of such mechanisms represents one side of the AI coin, but it is a side that frequently calls upon skills and techniques that fall outside of mainstream computer science: techniques from continuous mathematics, from optics, from information and control theory, as well as from various branches of engineering and materials science. In this second side of AI, which is quite different from the quasi-logical studies of knowledge representation, reasoning, computational complexity, and algorithm design, a great deal of help can be obtained from psychophysical and psychobiological studies of the capacities of the most highly developed intelligent systems we know, living ones.

We often cannot tell in advance whether certain problems belong in the first or second category of study. Thus for a long time it was generally accepted that perception was continuous with reasoning and involved recourse to prior knowledge, expectations, and inferences quite early in the process. Now researchers believe that a great deal of knowledge-independent processing goes on in so-called "early vision," and considerable progress has been made on this front by researchers familiar with optics, projective geometry, and neuroscience. The same is true of motor control. There has been considerable recent progress in understanding the human motor control system and properties of human muscle tissue, an understanding that is likely very soon to be exploited in the design of new robotic mechanisms.

The Fellows of the AI and Robotics (AIR) program have contributed to all aspects of the study of intelligent systems. The primary focus so far has been on "early vision" and on knowledge representation and reasoning. More recently we have developed exceptional strength in the basic science dealing with the design and analysis of autonomous agents — sometimes called "robotics" in its more applied guise. In these areas, the program has some of the best known scientists in the world, a conclusion that was clearly echoed by the members of the Review Panel which recently handed down its report on the first five years of the program. The program also has considerable strength in the contributing basic sciences, such as psychology, psychobiology, and mathematics. It has also begun to build additional strength on the architecture and motor control side, with the addition of robotics experts to our earlier complement of control theorists and neuroscientists specializing in human motor control.

Background: the Approach of the AI and Robotics Program of CIAR

Because AI is an *emerging* field, most of its concepts and techniques are very new (or at least have taken radically new directions in the past two decades) and developed as the field itself developed. Moreover, because many problems encountered are ones that have been studied by various disciplines, the pursuit of AI is essentially interdisciplinary. Research in AI requires cross-disciplinary skills in both content and methodology. In many cases large AI projects require both a degree of effort and a range of expertise that can only be achieved by interdisciplinary groups. However, because of the economic importance of the new technologies being developed under the heading of Artificial Intelligence, the competition for trained scientists is extreme. Because of this it is a field that is ideally situated to benefit from the kind of independent support of human resources as well as cross-disciplinary and cross-institutional interaction that CIAR can provide.

CIAR support is intended to help in the development of the science and technology of AI and cognitive science in Canada at three levels: in the personal development of promising researchers, in the development of a critical mass of researchers located at centers of excellence, and in the development of an interdisciplinary, geographically distributed research community across the country. CIAR concentrates its efforts on the fundamental science side of the basic-applied spectrum, but also contributes talent and advice to the more applied work focused within PRECARN and the various Provincial and Federal Centres of Excellence programs, such as IRIS, ITRC, and the British Columbia Advanced Systems Institute.

Background

There are said to be more than 50 AI centers around the world, the best of which are in the United States. However, Britain, Japan, France, Germany, and the European Economic Community have major national initiatives in this field, and other countries, like Holland, Sweden, and Australia also have national programs. Manpower in this field is extremely scarce, owing in part to its novelty and glamour, and to the fact that many large industries (particularly in the United States, Britain, and Japan) have recognized the economic potentials of the field and have mounted AI laboratories of their own — and are willing to spend a great deal of money to attract the few available AI experts. AI is a particularly appropriate field foı CIAR to support because it is an intellectually exciting frontier science which has great economic potential, because it is inherently an interdisciplinary field which can benefit from unconventional nondiscipline-based support, and because Canada already had a good base of internationally recognized experts who were, however, spread from Halifax to Victoria.

The AIR Program was the first program to be approved by the CIAR board. The approval occurred in 1983 and by the end of 1984 (the first full year of the program), an additional 14 people had been appointed: 11 Fellows and 3 Associates. Now there are 16 Fellows, 3 Scholars, 14 Associates (including 4 outside Canada).

In May of 1988 the program underwent a five-year review process by an international panel of experts and received high marks. Indeed, the committee placed the major centers (at McGill University, University of Toronto, and the University of British Columbia) among the top 10 in the world and attributed much of their success directly to the CIAR program. Fellows of the program have received many awards over the years (including a large number of "best paper awards," the IEEE Field Medal, the NRC Steacie Award, and the Hebb Award) and have held concurrent Fellowships such as the Killam Research Fellowship and the NRC Steacie Fellowship.

Program Activities

Meetings, Workshops, and Publications

The AIR Program has held a general meeting of all its members every year since its inception. In some cases, outside guests (or potential recruits) are invited as well. Attendance at these meetings is high. With a few exceptions everyone who can possibly make it does attend. It is at these annual meetings that many of the contacts are established which have led to further fruitful interactions.

In addition to the regular annual meetings, there have been a few special workshops put on by program members on various special topics. These include a workshop at Dalhousie in 1986 at which individuals presented, and even demonstrated, some of the research work being carried out by the Dalhousie neuroscience group and some other closely related work, one in 1986 on Computational Approaches to Problems in Human Vision, in 1987 on visually guided action, and one in 1988 on Vision and the Extrastriate Cortex. The proceedings of the first two workshops have now been published by Ablex as the first volumes of *The Canadian Institute for Advanced Research Series in Artificial Intelligence and Robotics.* Volume 1 is called *Computational Problems in Human Vision* and Volume 2 is called *Vision and Action: The Control of Grasping.*

Graduate Student Activities

Graduate students have also been part of our program activities. McGill University and University of Toronto vision research students (and more recently a couple of students from The University of Western Ontario) have participated in annual exchanges, with groups of 15 or so students traveling to the other laboratory for a two-day exchange of ideas and a discussion of their research. The interchanges among students working on vision was extended to a two-and-a-half-day workshop run by and for graduate students and held at York University. These graduate student meetings have now been formalized as annual events. The second one was held in the summer of 1988 in Edmonton, Alberta, on Knowledge Representation and Reasoning, the third was held in London, Ontario, in the summer of 1989 on Vision for Robotics, and a fourth is being planned for the summer of 1990 in the Vancouver area. In addition to these more formal exchanges and workshops, several students working with Program members were funded to travel to other CIAR Centers to discuss research topics with students at those nodes.

Table 3 gives some brief summaries of research interests provided by Fellows, Scholars, and Associates of the program in Artificial Intelligence and Robotics.

For additional information, please contact:

Zenon Pylyshyn, Director
CIAR Program in Artificial Intelligence and Robotics
Centre for Cognitive Science
The University of Western Ontario
London, Ontario
N6A 5C2
Telephone: (519) 661-2063

Supplementary References

Pylyshyn, Z. (Ed.). *Computational Processes in Human Vision: An Interdisciplinary Perspective.* Ablex Publishing Corporation, Norwood, NJ. 1988.

Goodale, M. (Ed.). *Vision and Action: The Control of Grasping.* ABlex Publishing Corporation, Norwood, NJ. 1990.

Table 3. Research interests of AIR Program members.

Members	Research interests
Fellows	
Zenon W. Pylyshyn, Director Centre for Cognitive Science Social Science Building The University of Western Ontario London, Ontario N6A 5C2 (519) 661-2063	Theoretical issues in human cognition; studies in human visual attention
Dr. Peter E. Caines Department of Electrical Engineering McGill University 3480 University Street Montreal, Quebec H3A 2A7 (514) 398-7129	Foundations of adaptive control, applications to machine learning
Professor Max Cynader Eye Care Clinic Vancouver General Hospital 2550 Willow Street Vancouver, British Columbia V5Z 3W1	Studies of the organization of the visual cortex, combining computational, anatomical, and physiological approaches
Professor Geoffrey Hinton Department of Computer Science University of Toronto 10 King's College Road Toronto, Ontario M5S 1A4	Neural network architectures and applications to learning, vision, and knowledge representation

Table 3. Continued.

Members	Research interests
Professor John Hollerbach Industrial Chair of Robotics McGill University Montreal, Quebec H3A 2A7	Theoretical studies in human motor control and design of intelligent robotic systems
Professor Martin D. Levine Department of Electrical Engineering McGill University 3480 University Street Montreal, Quebec H3A 2A7 (514) 398-7115	Design of general purpose vision systems; with applications in biomedicine and robotics
Professor Hector Levesque (Member Program Development Committee) Department of Computer Science University of Toronto 10 King's College Road Toronto, Ontario M5S 1A4 (416) 978-3618	Knowledge representation, models of belief, complexity issues in automated reasoning
Professor Alan K. Mackworth (Member Program Development Committee) Department of Computer Science University of British Columbia 6356 Agricultural Road Vancouver, British Columbia V6T 1W5 (604) 228-4893	Computational vision, representation of visual knowledge
Professor John Mylopoulos Department of Computer Science University of Toronto 10 King's College Road Toronto, Ontario M5S 1A4 (416) 978-5180	Knowledge representation, deductive data bases
Professor Raymond Reiter (Nodal Coordinator) Department of Computer Science University of Toronto 10 King's College Road Toronto, Ontario M5S 1A4 (416) 978-6324	Models of reasoning, nonmonotonic logics suitable for automated reasoning, application of logic to diagnosis, vision, and other domains
Professor Demetri Terzopoulos Department of Computer Science University of Toronto 10 King's College Road Toronto, Ontario M4S 1A4	Computational vision based on physical models; dynamic computer graphics

Table 3. Continued.

Members	Research interests
Professor John K. Tsotsos Department of Computer Science University of Toronto 10 King's College Road Toronto, Ontario M5S 1A4 (416) 978-3619	Biological and machine vision, complexity, and other theoretical issues in visual processing
Professor Robert J. Woodham (Nodal Coordinator) Department of Computer Science University of British Columbia 6356 Agricultural Road Vancouver, British Columbia V6T 1W5 (604) 228-4368	Computational vision (low-level), remote sensing
Professor George Zames Department of Electrical Engineering McGill University 3480 University Street Montreal, Quebec H3A 2A7 (514) 398-7141	Mathematical theory of systems with feedback
Professor Steven Zucker (Nodal Coordinator) (Member of Program Development Committee) Department of Electrical Engineering McGill University 3480 University Street Montreal, Quebec H3A 2A7 (514) 398-7134	Theoretical analysis of computer and biological vision
Institute Scholars	
Professor David Lowe Department of Computer Science University of British Columbia 6356 Agricultural Road Vancouver, British Columbia V6T 1W5 (604) 228-3170	Computational vision, object recognition
Professor Ian Hunter Biomedical Engineering Unit Faculty of Medicine McGill University 3775 University Street Montreal, Quebec H3A 2B4 (514) 398-6740	Theoretical and experimental studies in human motor control
Professor Mark Seidenberg Department of Psychology Stewart Biological Sciences Building McGill University Montreal, Quebec H3A 1B1	Investigations of language comprehension; models of human linguistic abilities

Table 3. Continued.

Members	Research interests
Professor Richard B. Stein Department of Physiology University of Alberta Edmonton, Alberta T6G 2E5 (403) 432-3796	Studies of human and animal motor control systems; design of motor prosthetic devices
Dr. William G. Tatton Playfair Neuroscience Unit Toronto Western Hospital North Pavillion, Room 903 399 Bathurst Street Toronto, Ontario M5T 2S8 (416) 369-5147	Neuroscience of motor control and visual-motor coordination
Foreign Associates	
Professor Wolfgang Bibel Institut fur Informatik Technische Universitat Munschen D-8000 Munchen 2 Germany	Automated reasoning, parallel inference machines, knowledge representation, program synthesis
Dr. Ronald Brachman, Head Artificial Intelligence Principles Research Department AT&T Bell Laboratories 600 Mountain Avenue Murray Hill, New Jersey 07974-2070 U.S.A. (201) 582-3000	Systems for knowledge representation and reasoning; frame-based description systems
Dr. Stanley J. Rosenschein Artificial Intelligence Centre ARI International 333 Ravenswood Avenue Menlo Park, California 94025 U.S.A. (415) 326-6200	
Associates	
Professor Roger Browse Department of Computer and Information Science Queen's University Kingston, Ontario K7L 3N6 (613) 545-6069	Studies of human and machine perception in visual and tactile modalities
Professor Mel Goodale Department of Psychology Social Sciences Building The University of Western Ontario London, Ontario N6A 5C2 (519) 661-2070	Visual control of action; neural basis of mammalian visuo-motor coordination

Table 3. Concluded.

Members	Research interests
Professor Alan Jepson Department of Computer Science University of Toronto Toronto, Ontario M5S 1A4 (416) 978-6488	Mathematical and computational models of biological vision systems
Professor Peter D. Lawrence Department of Electrical Engineering University of British Columbia 2356 Main Mall Vancouver, British Columbia V6T 1W5 (604) 228-5934	Design of robotic and telerobotic systems
Professor Susan Lederman Department of Psychology Queen's University Kingston, Ontario K7L 3N6 (613) 545-2878	Studies of haptic perception in humans, with applications to robotics
Professor Jeffry Pelletier Department of Philosophy University of Alberta 4-108 Humanities Centre Edmonton, Alberta T6G 2E5 (403) 432-5471	Issues in knowledge representation and natural language processing
Professor Len Schubert Department of Computing Science University of Alberta Edmonton, Alberta T6G 2E5 (403) 432-3330	Issues in knowledge representation and natural language processing. Design of language comprehension systems
Dr. Edward Stabler Department of Linguistics University of California, Los Angeles 405 Hilgard Avenue Los Angeles	Natural language processing and automated deduction
Professor Anne Treisman Department of Psychology University of California Berkeley, California 94720 (415) 643-8408	Studies of human attention; visual processing

Centre de recherche informatique de Montréal

Keywords: Collaborative research, artificial intelligence, expert systems, knowledge acquisition, speech, neural networks, user interfaces, qualitative reasoning, diagnostic systems, software engineering, robotics, simulation, deep knowledge, engineering, meteorology, business, natural language, telecommunications, technology transfer, seminars, courses, workshops

Introduction

The Centre de recherche informatique de Montréal (CRIM) is a nonprofit corporation registered under the Québec Companies Act, composed of 7 university members, 30 industrial members, and 57 employees. The official inauguration of CRIM in October 1985 concluded the efforts of a number of people working in universities, industries, and the government to set up, in Montreal, a computer research center where researchers from the university and industrial sectors could collaborate.

CRIM has designated artificial intelligence as a research priority. The other aspects of computer research at CRIM reflect the central theme; often, several research domains are involved in a single project. These domains include software engineering, man–machine interfaces, parallel computing, computer communications, CAD/CAM, and VLSI. CRIM's laboratories offer a favorable environment for both research and development projects.

Expert systems

Expert systems and knowledge-based systems are a major area of research and development at CRIM. This technology is applied to various domains, in different projects.

Research

While implementing a number of different applications, CRIM has continued to pursue fundamental AI research. The areas in which research continues to be carried out are:

- development and study of knowledge acquisition techniques and models
- automation of the knowledge engineering process
- incorporation of deep knowledge in expert systems
- automatic testing and validation of expert systems
- using very large knowledge bases and complicated causal networks
- development of automatic speech recognition
- application of neural networks to phoneme classification

CRIM has made significant advances in:

- blackboard systems
- integration of statistical, heuristic, model-based, and case-based reasoning strategies
- development of intelligent user interfaces
- qualitative reasoning techniques for time-varying systems
- integration of qualitative reasoning and numerical optimization
- constraint satisfaction algorithms for qualitative reasoning

Hydroelectric Applications

A number of projects involve the hydroelectric industry. A "family" of expert systems was proposed to perform diagnoses on a set of electrical equipment, in terms of both maintenance and repair. The project involved developing a knowledge acquisition model for the design of a general diagnostic expert system for different types of electrical equipment.

In conjunction with several of its members, CRIM takes part in a project which involves research into numerical methods for hydraulic installations. The following subjects are being investigated: databases related to engineering, digital cartography, and knowledge-based systems in engineering design. The main goal is to automate the various phases of knowledge engineering and to optimize certain tasks on computer-assisted design products.

The development of an expert diagnostic system for monitoring the state of hydroelectric dams is a project where CRIM is presently in the modeling phase. The model is intended to be sufficiently flexible and generic to be adaptable to other different work.

A project where a decision-making aid for optimizing software for simulating hydroelectric complexes was the main topic is now complete.

An intelligent robot is the subject of a joint research project in which CRIM and two other research centers are participating. The robot will be guided remotely by an operator and will automate the maintenance of high tension lines. The robot will have its own knowledge base which will permit it to perform certain tasks automatically and to consult the operator on its decisions.

Process Industry Applications

The process control industry is also benefitting from CRIM's research and development expertise. The solving of problems caused by pitch deposits in kraft mills is being addressed in a project based on the expertise of a human pitch expert. Deep knowledge of the kraft process is incorporated in the system.

CRIM has completed a feasibility study with one of its industrial members for a project concerning a diagnostic expert system for the metallurgical industry. Research topics for the project may include methods of dealing with a very large knowledge base and reasoning with a complicated causal network.

Meteorological Applications

A prototype expert advisory system to assist meteorologists in the production of airport terminal weather forecasts is being built. This system will simulate qualitatively the physics of certain phenomena that are not included in the numerical models which the meteorologist currently uses for his predictions. In addition to meteorological research, there is on-going research in methodologies for the testing and validation of expert systems.

General Business Applications

With regard to more general business applications, the modeling of corporate data was the focus of an early project at CRIM. A prototype data modeling expert system was developed.

CRIM has also developed an expert system prototype which diagnoses an enterprise's condition by means of the analysis of a series of objective and subjective variables.

Other Applications

In the area of speech recognition, CRIM has created a vocal interface for a medical expert system. Currently, a printed circuit board for speech recognition in both English and French is being developed. CRIM is also creating a system to recognize and search for specific keywords in continuous speech (word spotting).

The analysis of natural language dialogues and the generation of text are on-going research challenges at CRIM. A database interrogation tool which permitted bibliographic queries was designed and developed.

In the engineering domain, a knowledge-based system to assist in the verification and troubleshooting of flight simulation models is the object of another project.

With regard to software engineering, a study of development tools and methodologies for software incorporating artificial intelligence concepts is being developed by one of CRIM's researchers.

In the area of telecommunications, CRIM is currently developing protocol implementations in nonprocedural languages. The definitions contained in the protocol specification will be translated into rules which are used as input to a fifth-generation interface engine.

Laboratories

Since their inauguration on April 25, 1988, CRIM's laboratories have become a showcase for state-of-the-art computer science. The laboratories thus have both responded to, and increased the interest of, the manufacturing and service sectors in advanced computer technologies.

The laboratories are an important component of CRIM's services, playing three principal roles. The first is to transfer technology through seminars, courses, and workshops. These range from a general overview of a topic to the use of specific application software. The second is to offer members and nonmembers an opportunity for evaluating hardware and software, so that they can make informed choices. The third is to supply users with advanced facilities for project realization and research.

In order to be attentive to the needs of the Quebec community, CRIM's laboratories contain a wide range of systems, from PCs to super-mini computers. The available software is representative of the tools available in R&D laboratories as well as in advanced computer companies throughout the world. Specific application software available to users includes both advanced CASE tools and state-of-the-art expert system shells.

For further information, contact:

Dr. R. De Mori
Vice-President, Research and Development
Centre de recherche informatique de Montréal
1550 Maisonneuve Boulevard West
Room 1000
Montreal, Quebec
H3G 1N2

Information Technology Research Centre

The Information Technology Research Centre (ITRC) is a center of excellence supported by the Ontario Premier's Council Technology Fund that promotes fundamental and applied research in information technology at the Universities of Toronto, Waterloo, Queen's, and Western Ontario. Artificial intelligence, cognitive science, and intelligent control, as a group, form one of ITRC's five research areas. In this area, the research ITRC supports at the four universities focuses on the following topics.

Knowledge-based Management Systems

Professor Mylopoulos at the University of Toronto is leading a project to develop a system that facilitates the creation, access, and management of large knowledge bases. The main thrust of the project is to establish that efficient and robust implementations of such systems are possible through the adoption of techniques from database management for concurrency control, error recovery, and query optimization. A knowledge representation language, Telos, has been developed for experimenting with the implementation of knowledge-base management systems.

A second thrust of the project is to adopt recent results from natural language and knowledge representation research to develop more friendly user interfaces to a KBMS. A final, more theoretical, thrust of the project is to make advances in areas of knowledge representation, natural language, concurrency control, and query optimization that are relevant to the prototype system under development.

Knowledge Acquisition, Synthesis, and Organization

Professor Wong at the University of Waterloo is leading a project to apply pragmatic techniques for knowledge acquisition, including inductive learning based on probabilistic inference, clustering of data or knowledge, and pattern discovery, in combination with knowledge synthesis techniques for application to image databases. These techniques are robust in the presence of noisy data or uncertainty, and several companies have collaborated on their use in machine vision and vision knowledge systems.

Recently a new approach to 3-D model representation based on hypergraphs and rule network modeling capable of recognizing and locating fairly complex 3-D objects placed in unconstrained and cluttered environments has been developed.

Control of Discrete-Event Systems

Discrete-Event Systems (DES) arise in service and logistic systems, manufacturing, and communications, as well as in many other domains such as vehicular traffic and robotic and process control at the level of task coordination.

The increasing complexity of man-made DES, made possible by computer technology, underlines the need for formal theories of and systematic design approaches to DES control. The project "Control of Discrete-Event Systems" led by Professor W.M. Wonham (Toronto) with Professor G.K.H. Pang (Waterloo) is directed to these objectives. Research in progress focuses on how observations of the behavior of a DES can be exploited for its control and how controls for large complex DES can be built up systematically from simpler modular components.

Partners in Ontario industry have stimulated specific direction of project efforts to the study of communication protocols and manufacturing information systems.

Biological and Computational Vision

One objective of this project at Toronto, led by Professor Tsotsos with Professors Jepson and Milios, is to study object recognition in the presence of sensor noise, object occlusion, and object deformation. A contour-based approach is used to generate a small number of candidate hypotheses and to describe deformation, while a spatial occupancy represention is used to select the correct hypothesis. If noise is present, it is difficult to reliably trace contours. In that case, a method based on tangent fields and the Hough transform is being explored. Finally, extensions to the methods of the matching of three-dimensional range images are being investigated. These methods will be useful in industrial inspection.

Another objective is to determine the mechanism by which selective visual attention is implemented in biology. Since there is a large body of experimental data with which to compare, the researchers are striving to achieve the same visual performance that is observed in primates and in humans. Hopefully, this will lead to functional computational systems incorporating attentional capabilities, something that is missing from current computer vision models. The key idea is that attention must affect virtually all levels of visual processing and not only the so-called "cognitive" levels as previously believed.

A third objective is to determine an appropriate representation for the several levels of transformations that motion information undergoes as it is processed, with emphasis on a biologically plausible representation. The motion hierarchy will be only a portion of the entire intermediate level vision process and must be integrated in the future with shape, stereo, and color, for example. The biological plausibility constraint is important because it also includes the requirement of efficient computation by small spatially parallel computation structures.

The project team is attempting to develop a robot which uses active vision to move autonomously from one position to another within a large industrial bay. At all times it will know where it is on its map and as it moves it will avoid any objects, either moving or stationary. It will be able to recognize objects, such as valves or gauges, whose shape and size, but not location, are stored in an internal database. It will also be able to produce a contour map of any variable, such as noise level, that it can measure with its sensors in any designated area within the industrial bay. Active vision refers to a system that provides scene information by actively changing the geometry of, or moving, the optical system, not to the use of structured lighting systems or laser range finders. An active vision system would be able to change the focal length of a camera, change the relative positions of cameras in a binocular system, or rotate or move the cameras. The second objective of the project is to equip the robot with the ability to form a map of its own environment, by exploring it using its vision sensors.

Learning Procedures for Massively Parallel Networks

Professor Hinton is leading a project at Toronto to develop neural network learning procedures that are more efficient, more widely applicable to real tasks, and easier to implement in biological hardware or on a silicon chip.

He has recently developed an unsupervised learning procedure that discovers underlying features which are invariant across time or space. This procedure can be applied to a wide variety of perceptual tasks. It should work much faster than existing procedures in large networks because it allows many separate modules to learn in parallel with no interference.

Professor Hinton also developed an efficient version of the Boltzmann machine-learning procedure for recurrent networks of analog units. This procedure is very simple to implement in analog hardware. To test the usefulness of these new learning procedures they have been applied to the task of speech recognition. This is an important application area in which many different techniques have already been tried. Time-delay neural networks are now the best existing technique for the discrimination of very similar consonants. We are developing new kinds of neural nets that can deal effectively with variations in speaking rate, and with changes in phonetic context, loudness, and speaker identity.

Another important application area is shape recognition. Professor Hinton has developed an efficient way of encoding spatial relationships in neural networks. This allows a small module to recognize a particular object from any viewpoint. The invariance learning method is currently being applied to allow a self-organizing network to create these object modules automatically.

Finally, a new way of using neural networks to allow humans to control complex physical devices has been demonstrated. Using a "dataglove" a person provides input to a neural network that then produces output which drives a speech synthesizer.

Experimental and Computational Studies of Visual and Tactile Perceptual Processes

Professor Browse leads an interdisciplinary team at Queen's University, comprising Professors Lederman, Ellis, and Blostein, to study several aspects of visual and tactile perception.

Computational Studies of the Selectional Processes of Perception

Most computational vision systems make the assumption of the availability of clear well-centered uniform resolution images of the class of objects in question. The reality of human vision is that perceptual processes of varying resolution must be aligned with objects of interest during the process of perception. Perceptual systems for autonomous robots will face the same difficulty. In this project systems have been developed which intelligently select processing locations within an image utilizing a graded resolution input of similar structure to the human retina. These systems select locations on the basis of image information content, the status of ongoing interpretation, and on the basis of task requirements.

Studies in Haptic Attention

This project investigates some fundamental characteristics of human haptic perception. While numerous studies suggest the nature of appropriate psychophysical units for visual perception, relatively little is known about human haptics. The studies proposed continue the investigation of the nature and organization of feature properties which may be utilized in haptic perception. In addition, a related program has begun to examine many other aspects of haptic perception using a new apparatus, Q'Hand, that has been custom designed for this purpose. Q'Hand automates the preparation of stimulus displays, and their presentation to the fingertips.

As the presentation of these tactile displays is also automated, more than eight times as much data can be collected as has been possible with manual displays used in the past. As an important consequence, Q'Hand achieves a marked increase in the precision of tactile data that may be collected; moreover, it permits the researchers to pose questions concerning haptic attention and recognition that were previously difficult, if not impossible, to investigate.

Segmentation, Surface and Object Recognition

The human visual system is structured to provide limited perceptual processing in parallel across each retinal location. These processes are responsible for image segmentation and edge detection. Recent experiments with human subjects have focused on the determination of the extent of these parallel processes in order to gain the insight necessary to develop general-purpose computational vision systems. The results indicate that the human processing capabilities in early preattentive segmentation are best described as constraints upon banks of image filters rather than as properties of image features. Computational segmentation systems based on this principle are currently being developed. Other experimentation has shown that the boundary between parallel and serial processing in human tactile perception exhibits many of the same characteristics as are found for vision.

Object recognition requires the availability of models which depict the objects in question. The choice of representational strategy must be made in conjunction with the development of processes which extract information from the image. In this project, variations on Beiderman's Recognition by Components System have been developed. This approach specifies object models as combinations of a small fixed set of geometric primitives. The Queen's teams' extensions include using multiple resolutions and more sophisticated matching algorithms.

Experimental and Theoretical Studies in Cognitive Science

Professor Pylyshyn is leading a team including Professors Campbell, Goodale, Gurnsey, and Humphrey at The University of Western Ontario, which is directed at understanding some of the features of human visual attention and human visual processing where we are in some ways extremely facile (e.g., tracking moving objects, perceiving the shape and 3-D motion of objects from texture and 2-D motion fields; coordinating reach and grasp actions with visual information) and also those features where we appear to exhibit limitations (e.g., keeping track of several independent objects within a field of view that contains distracters with many of the same features; doing mental arithmetic and other tasks which stress working memory capacity; detecting certain kinds of changes over others, e.g., color change vs. brightness change).

Experimental evidence of these perceptual-cognitive abilities, along with computational models of them and computational tools for exploring such issues, forms an important basis for the design of automated systems — from telerobotics to human–computer interfaces. The work we are doing — both fundamental and applied — is a contribution to understanding the strengths and limitations of human information processing, and therefore of the foundations for the relationship of human to machine skills.

Application of Image Analysis to Radar Storm Tracking

In this project, Professors Pylyshyn, Mercer, and Barron at The University of Western Ontario are applying image analysis and artificial intelligence techniques in problems of weather forecasting. Of particular importance is the tracking of southern Ontario storms in radar data produced at the radar station at King City, Ontario. A particularly well-established AI technology, pyramid data structures, is being used to locate storm centers in the radar data. Further investigation of other image analysis techniques is underway.

Future work will generate an expert system to track storms through time for the meteorologists at King City. A prototype expert system has been developed and work continues on the knowledge acquisition phase of the expert system.

Motion Analysis for a Machine Vision System

"Motion Understanding" in the field of computer vision requires that sequences of two-dimensional images be analysed for changes that result from the relative 3-D motions of objects with respect to the observer. A project team including Professors Pylyshyn and Barron has designed a set of spacio-temporal filters that measure the various image deformations such as translation, curl, vorticity, and shear that are occurring in a fixated image of some surface over time. The motion of the surface as well as its local structure can be recovered from these deformations. Despite the theoretical success of our formulation for this problem, the recovery of motion and structure may be too sensitive to noise in the input image data to be of practical use. However, progress is being made in this area.

For further information, please contact:

Vic DiCiccio, Associate Director
Information Technology Research Centre (ITRC)
ITRC Site Office, Davis Centre
University of Waterloo
Waterloo, Ontario
N2L 3G1
Telephone: (519) 888-4540
Fax: (519) 885-1208
E-mail: vfdiciccio@ccng.uwaterloo.ca

PRECARN

Keywords: Research, industrial consortium, precompetitive, university–industry linkages, robotics, knowledge-based systems.

Abstract

PRECARN is an industrial research consortium which has been created to carry on and promote long-term, precompetitive research and experimental development in the fields of advanced robotics and artificial intelligence. Its membership includes over 30 Canadian corporations drawn from a broad cross-section of economic activities across Canada. It operates on the strength of annual and project-specific contributions from its member companies and utilizes existing government support programs for its research efforts. This paper reviews the motivation, objectives, and accomplishments of this unique research consortium.

Introduction

The name PRECARN is an acronym for "PRE-Competitive Applied Research Network." It is an industrial research consortium created to carry on and promote long-term applied (or precompetitive) research in the fields of artificial intelligence and robotics.

Motivation

The common thread that brings and holds the diverse membership together is an appreciation of the increasing importance of robotics and artificial intelligence technologies in a rapidly changing global economy. All member corporations, whether they be potential users or producers, recognize the increasing impact of these technologies on our collective ability to remain competitive as a trading nation. They see the need, therefore, to be actively aware of, and involved in, the total spectrum of research leading from fundamental knowledge to the completion of new product development.

PRECARN's members recognize another reality. That is, that no one corporation, regardless of its size, can on its own be completely effective in the total spectrum of research. This is especially the case with regard to precompetitive research, that band of activity where one is working with new fundamental knowledge, but without a full appreciation as yet of its shorter- or longer-term applications. It is the need for collective action in this precompetitive area that has brought PRECARN together, just as it has prompted the formation of consortia worldwide. Many of these other consortia are industry specific; PRECARN, however, is pursuing a technology-specific program in recognition of the broad, industry-wide applications of artificial intelligence and robotic systems. This collective approach will ensure not only more efficient use of research funding from all sources, but also more efficient use of the thinly spread research community itself.

Objectives

PRECARN's overriding objective is the creation of greater economic wealth in Canada through increased industrial competitiveness. To do this, we must create a greater "receptor capacity" in industry for new knowledge on, and new applications of, artificial intelligence and robotics. In fact, it was the realisation that Canadian industries needed to strengthen this receptor capacity which prompted the Canadian Institute for Advanced Research (CIAR) to initiate the actions which eventually led to the creation of PRECARN in 1987 by a founding group of 22 Canadian corporations.

PRECARN's research program is the main instrument for helping develop this strong receptor capacity and meet the following objectives:

- To bring private sector users and producers of this technology together.
- To develop effective and lasting university–industry cooperation.
- To identify long-term industry needs, challenges, and opportunities.

Additionally, PRECARN aims at developing a highly efficient "exploitation mechanism" to get the product of research out to members. Its policy on intellectual property ownership encourages this by providing royalty-free licensing rights to all members for technology resulting from PRECARN-funded projects.

Lastly, to fund this research program, PRECARN aims at leveraging its private sector funding to maximize government contributions from existing or new research granting programs.

Organization and Structure

PRECARN's current membership consists of 34 Canadian corporations drawn from a broad cross-section of economic activities across Canada. These include:

- Mining and Mineral Processing
- Oil and Gas
- Electric Utilities
- Telecommunications
- Aerospace
- Electronics/Computer
- Federal and Provincial Research Institutes

The membership includes users as well as producers of AI/robotics technology. A list of current members is provided later in this report.

PRECARN's management structure consists of a Board of Directors, an Executive Committee, a President/Chief Executive Officer, and a Vice-President. They are supported by a small staff and four Advisory Committees. It was the wish of the founding members that PRECARN be managed at the Board of Directors level with a very high representation from the membership. For this reason the number of Directors was set at 25, and they are nominated from the membership's list on a rotation basis.

It was also the wish of the members that PRECARN be a research institute "without walls," similar in that respect to the CIAR. PRECARN's research projects will typically be carried out jointly between member organizations and university and/or government research laboratories. This allows PRECARN to operate with a small staff and thus maximize the use of its funds for the actual research.

PRECARN's four advisory committees represent the key mechanism through which the corporation will achieve its goals of coordination, evaluation, dissemination, and exploitation of the research. These committees are appointed by the Board of Directors primarily from representatives of the member companies but also, in the case of the Technical Advisory Committee, from the broader Canadian research community. These committees, which report to the Board of Directors, have the following roles:

- *Policy and Planning*: Responsible for preparing and updating the long-range corporate plans.

- *Audit and Evaluation*: Responsible for arranging proper financial audits and reports and also a comprehensive audit on the management of the total system.

- *Exploitation*: This committee might be looked upon as the marketing department of a conventional corporation. It establishes and monitors progress on the various processes for the early identification of exploitable results.

- *Technical*: This is the largest committee as it represents in many ways the heart of PRECARN's operations. In short, its role is to design, recommend, review, and evaluate the research program.

The Research Program

PRECARN's intent is to develop a major artificial intelligence and robotics research program in application areas that hold the greatest promise for eventual exploitation by its member companies. This research program is expected to draw on the combined strengths of university, government, and industrial research laboratories and enhance linkages among them. PRECARN's primary interest is in long-term, precompetitive research where the results in terms of knowledge, processes, etc., are of generic interest and will need further work to adapt or develop them for specific applications.

The process by which PRECARN identifies and selects the research projects that will form part of the overall program is a two-step competitive operation. First a Request for Proposals is issued, inviting applicants to submit brief outlines of research projects that meet PRECARN's objectives along with a detailed Feasibility Study Proposal. Successful applications are then funded for the feasibility study phase to a maximum of $100 000. The specific number of Feasibility Studies accepted depends on the adequacy and relevance of the proposals as well as available funding. They are expected to be completed within six to nine months and to include the preparation of an expanded research proposal. Mandatory criteria for projects include the following:

- Must be in artificial intelligence or robotics
- Must have potential for creation of new wealth in Canada
- Must contribute to Canada's industrial capacity to receive, understand, and apply new AI/robotics technologies
- Must be longer term

- Must have at least two PRECARN members as sponsors
- Should develop effective and lasting university–industry linkages

The second stage of the process is to select from the completed feasibility studies the research projects that PRECARN feels hold the greatest potential for ultimate benefit to its members. PRECARN then proceeds to develop a financial support package built around its own funds, funding from PRECARN members with particular interests in the project, and government programs. Once approved, the research projects are expected to last between three and five years, during which time there will be regular briefing sessions and workshops open to all PRECARN members on the project's progress and results. While ownership of any intellectual property generated during a project will be negotiated on a case-by-case basis at the outset of the project, in all cases, PRECARN's full membership will have a royalty-free right to license.

Accomplishments to Date

The first and maybe the most important accomplishment has been to bring together 34 companies that have the vision and the commitment to embrace the PRECARN objectives. Each company's annual contribution is $25 000; in addition they contribute time, effort, and costs in involvement on the Board of Directors and the four standing committees. And of course they contribute through their involvement in the research projects themselves.

The second accomplishment has been the negotiation of a Membership Agreement that deals with sensitive matters such as the ownership of intellectual property. This agreement has been reviewed and accepted by 34 corporate legal departments and positively reviewed by the Competition Policy Branch of the federal government.

Third, PRECARN's Research Program was initiated in January 1988 through a Request for Research Proposals. Twenty-eight proposals were received and a first group of seven selected for feasibility funding in April 1988. Two others have subsequently been approved, bringing the total investment by PRECARN in feasibility studies to $884 000. Six studies were completed in June 1989, and in September the Board of Directors selected four of them for the research program, planned to get under way in early 1990. It is also expected that additional projects selected from the remaining feasibility studies will be initiated later in 1990 and would bring the total research effort to more than $7 million per year over the next five years. PRECARN is already planning for the next rounds of Request for Research Proposals and Feasibility studies. In a few years the total PRECARN annual research program budget could reach $15 million.

Fourth is the funding commitment which PRECARN has obtained so far from industry and government. Through their membership fees and in-kind contributions, the PRECARN members have already invested close to $4 million in this new consortium. PRECARN will continue to provide the feasibility program, overall management, and a project contingency fund. Individual members will also contribute a minimum of 25% of the cost of each research project. Federal and provincial governments are contributing to PRECARN's Research Program through the participation of a number of government research laboratories in the projects and through funding contributions to the overall program. Currently in place is a contribution arrangement with the National Research Council, under their Industrial Research Assistance Program (IRAP), which has already provided some funding assistance for feasibility studies and for project administration. Total assistance via this route could reach up to $1 million per year. Another contribution agreement has been concluded with Industry, Science and Technology Canada, under the Strategic Technologies Program, for funding assistance of up to $10 million over five years. In addition PRECARN is still seeking provincial support which, most likely, will be in two areas; financial assistance to small- and medium-sized companies to join PRECARN and project-specific support. Already Alberta and British Columbia have approved the former and will contribute up to 50% of the PRECARN membership cost of eligible companies.

A fifth accomplishment has been the establishment of solid linkages with the Canadian university research community. This is probably best evidenced by the Institute for Robotics and Intelligent Systems (IRIS), a nation-wide research program in intelligent systems, managed by PRECARN and involving over 120 academic researchers, including Canada's top talent in AI and robotics. IRIS is one of 14 selected proposals under the Federal Networks of Centres of Excellence Program.

Conclusion

PRECARN is a young research organization, based on the concept that an industry-led consortia emphasizing linkages between university, government laboratories, and industry will provide industry with the necessary knowledge and tools to remain competitive in today's highly technological and global economy.

Many see PRECARN as an experiment, and while it is true that PRECARN has had to blaze new paths almost every step of the way, the results so far are extremely encouraging. Industry and government have, by their contributions to-date to the program, demonstrated their support of the concept. The essential university–industry linkages have been demonstrated by the active participation of university researchers in all our projects and with the IRIS proposal.

The network and linkages have been established, the initial funding put in place, and the initial research investigations completed. PRECARN now enters the main phase of its development: the research program itself.

PRECARN Members

- Alberta Research Council
- Alcan International Limited
- Atomic Energy of Canada Limited
- British Columbia Advanced Systems Institute (BCASI)
- British Columbia Hydro
- Bell-Northern Research
- Bristol Aerospace Limited
- CAE Electronics Ltd.
- Canadian Institute for Advanced Research (CIAR)
- Canadian Marconi Company Limited
- Centre de recherche informatique de Montréal (CRIM)
- EBCO Industries Limited
- Falconbridge Limited
- Hatch Associates Limited
- W.G. Hutchison & Company Limited
- Hydro-Québec
- Inco Limited
- Lac Minerals Limited
- MacDonald Dettwiler and Associates (MDA)
- Manalta Coal Limited
- Microtel Pacific Research Limited
- MPB Technologies Inc.
- National Research Council of Canada
- New Brunswick Power Commission
- Noranda Inc.
- Oerlikon Aerospace Limited
- Ontario Hydro
- ORTECH International
- Petro-Canada Resources
- Shell Canada
- H.A. Simons Limited
- Spar Aerospace Limited
- STELCO Inc.
- TransAlta Utilities Corporation
- TransCanada Pipelines
- Xerox Research Centre of Canada

List of Projects

- Telerobotic Development System
  MPB/Hydro-Québec/CAE/CRIM/McGill University/CIAR
  Approved for Research Phase

- Mobile Robotic Platform for Rugged Terrain
  MDA/BCASI/University of British Columbia
  Feasibility Study Completed

- Active Navigation of a Mobile Robot in a Known Environment
  Ontario Hydro/AECL/NRC/University of Toronto/York University
  Approved for Research Phase

- Intelligent Robotic Control for Mine Drilling and Ore Handling
  Inco/Falconbridge/McGill University/École Polytechnique de Montréal
  Approved for Feasibility Funding

- An Advanced Real-Time Expert System
  MPB/Noranda/NRC/University of British Columbia/University of Waterloo/Queen's University/University of Ottawa/University of New Brunswick
  Approved for Research Phase

- Applications of Expert Systems to Advanced Process Control
  Ontario Hydro/CAE/University of Toronto
  Approved for Research Phase

- An Environment for the Development of Integrated Intelligent Forecast Maintenance Systems

Hydro-Québec/Alcan/CRIM/McGill University
Feasibility Study Completed

- Application of AI to Engineering Design
  Spar/Bristol/Marconi
  Feasibility Study Completed

- An Intelligent Graphic Interface for Real-Time Monitoring and Control
  MPB/BCASI/Simon Fraser University
  Feasibility Study Completed

For additional information, contact:

G.M. MacNabb, President and CEO
PRECARN Associates Inc.
30 Colonnade Road, Suite 300
Nepean, Ontario
K2E 7P4
Telephone: (613) 727-9576
Fax: (613) 727-5672

PRECARN'S Institute for Robotics and Intelligent Systems Network

The Institute for Robotics and Intelligent Systems (IRIS) is one of 14 research networks recently selected by the Canadian Federal Government under the new $240 million Networks of Centres of Excellence Program.

IRIS is a collaborative effort between the academic research community and the private sector. It brings together within one nation-wide network the leading academic talents in the fields of artificial intelligence and advanced robotics, including the fellows of the Canadian Institute for Advanced Research. IRIS is managed by PRECARN, an industrial consortium for precompetitive research in AI and robotics, and is operated as a separate component to PRECARN's own industrial research network, but with the same procedures for project review and knowledge dissemination.

The research program of IRIS includes 22 projects within three areas of investigation: computational perception, knowledge-based systems, and intelligent robotic devices. The program will be carried out over four years, with a federal grant of $23 million. Over 100 researchers from across Canada will participate.

The following description of the overall research program and its 22 projects is reprinted from the original IRIS submission.

The Scope of the Research Program

The research program envisaged for IRIS encompasses three major areas: Area A, Computational Perception; Area B, Knowledge-based Systems; and Area C, Intelligent Robotic Systems. Intelligent Robotic Systems in this instance includes any system where the outputs of multidimensional sensors are fused in real time through computational means and then used by an internal knowledge base, and supplemented by any desired level of human supervision, to energize actuators which cause a machine or system to carry out a desired task. This definition therefore includes a broad range of expert and robotic systems, but clearly one of the long-term objectives is to develop intelligent robotic devices that can work in unstructured environments and very hostile conditions such as those presented by our resource industries. Advances towards this eventual objective will have major impacts upon the efficiency and therefore the competitiveness of the Canadian manufacturing, transportation, service, and resource industries as outlined in a later part of this presentation.

Area A: Computational Perception

This area is driven by the requirement that an intelligent robot must be able to perceive its environment. Research in computational perception involves the embodiment of perceptual theories in the design and implementation of systems. These implementations, in software and hardware, also serve as the basis for development of technology for automation, telerobotics, remote sensing, and biomedical applications.

The main research topics proposed are computational vision (Projects A1–A6), tactile perception (A2, A4), 3-D sensing (A4, A5), auditory perception (A3, A7), and sensory integration (A2–A6). The challenge in computational vision is to infer scene descriptions from images in various forms. Applications in robotics, telerobotics, and remote sensing require solutions to this problem. The research tasks involve the use of knowledge of shading, contour, stereo, scale, motion, color, and surface and object models to extract useful scene descriptions. There are mathematical, physical, computational, and architectural problems to be solved in order to develop high-performance reliable systems for these applications. For example, the mathematics of differential geometry, the physics of image formation, shape-matching algorithms, and the use of new sensor designs and parallel architectures will all be investigated. This research will enable the development of applications such as a model-based system for visual feedback in telerobotics (A1) and navigation systems for mobile robots (A6). These applications require new object and spatial layout representations, active vision, new sensor and processor designs (A1, A2, A4, A5), and sensory integration: the fusion of information from different sensor types such as TV cameras, laser rangefinders, and tactile force sensors.

Area B: Knowledge-based Systems

Breakthroughs in hardware and software technologies have made it possible to build systems that exhibit "intelligence" in the performance of a task, ranging from medical diagnosis to design of computer installations to interpretation of geological data. The nucleus of such expert systems consists of a knowledge base storing knowledge about the application domain (medicine, computer systems engineering, or geology) and an inference engine which generates inferences from a set of input data using the contents of the knowledge base. Considerable research has been generated in recent years on the formal properties of knowledge, the formal characterization of reasoning patterns, the data structures, and algorithms needed to support knowledge-base management.

The main research topics are the following:

- *Knowledge-base Management*: How to construct, access, and maintain a large knowledge base, how to optimize queries, how to offer the knowledge base to concurrent users (Project B3).
- *Reasoning*: Studies of different types of reasoning, such as nonmonotonic, probabilistic, and abductive reasoning. Also, the development of efficient implementations of inference mechanisms for logic programming and constraint logic programming (Projects B5 and B6).
- *Human Interfaces*: How to build various types of interfaces between a knowledge base and its human users, including natural language front ends, speech and text processing interfaces, and graphical interfaces (Project B2).
- *Applications of Knowledge Bases to the Design of Intelligent Information Systems*: This project will explore how to facilitate the development of information systems and how to build more effective systems by using knowledge representation (new design languages), knowledge bases (expert systems for design), and also by viewing an information system as a knowledge-based system that reasons with respect to what it "knows" in accepting new information and in responding to questions (Project B1).

Area C: Intelligent Robotic Systems

Research in this area addresses a wide range of challenges that currently inhibit the use of robotic devices for anything more than preprogrammed repetitive actions in relatively benign surroundings. The nine research projects have been selected by IRIS and PRECARN for this area have the collective objective of drawing on, and adding to, the research in Areas A and B so as to arrive eventually at a point where, first, telerobotic devices and, eventually, autonomous robotic devices, can work effectively in even the most hostile environments such as in space, underwater, and in our mines and forests. This ultimate objective of intelligent robotic systems is well beyond the four-year timeframe of this research submission, but the projects that are proposed will move us a considerable distance toward that objective by building prototypes and addressing the more immediate challenges.

Area A: Computational Perception

A1: Vision Systems for Recognition, Tracking, and Navigation

Project Leader: Dr. A.K. Mackworth

The goal of this project is to build reliable and efficient vision systems for object recognition, localization, inspection, and tracking and for navigation. This requires development of robust algorithms for early vision implemented on special-purpose hardware, new intermediate representations and algorithms for depth and orientation maps, and new object representations and shape matching algorithms for articulated, curved, and flexible objects. These systems will be used in robotics, telerobotics, and diagram understanding applications.

A2: Human and Machine Perception Research for Advanced Telerobotics

Project Leader: Dr. S. Lederman

Telerobots are human-operated robotic manipulators typically used in dangerous environments such as undersea, space, and in nuclear reactors. Humans are included as part of the function of these devices largely because the sciences for robotics and artificial intelligence are not sufficiently advanced to permit the development of autonomous robots for these applications. To permit the operation of these telerobots in a wider variety of situations including complex and delicate tasks, it is necessary to develop computational perception support systems, and it is also necessary to understand the human's ability to cope with information fed back from the robot site.

Members of the Queen's University Robotics and Perception Laboratory have expertise in computational aspects of both tactile and visual perception, robotics including control and force reflectance, human factors, and the psychology and psychophysics of human visual and tactile perception. These are the key ingredients required for a coordinated research project focused on the perceptual needs for telerobotic operation. Within this project we will build a common facility with which to experiment in force reflectance, tactile and visual perception support systems, intersensory integration, graphic display of computational results, and human factors.

A3: Architectures for Sensory Processing

Project Leader: Dr. M.S. Cynader

Intelligent robotic systems will have to be able to combine sensory inputs from a variety of sources to effectively sense their environment. The raw sensory image contains information that must be processed and reconfigured to allow inferences about the trajectories and locations of both objects and surfaces in three-dimensional space. The research proposed here is novel and benefits in a unique way from the network in which it is embedded. By linking the biological vision and neural modeling expertise of the Cynader group with the computational strengths of other nodes within the network, we expect to achieve an exciting synergy between different approaches to common problems. Evidence of this synergy is already appearing in the recent paper in *Nature* involving Cynader's group and Steven Zucker, who is a key participant in the McGill University node of the IRIS proposal.

A4: Multisensor Perception

Project Leader: Dr. M.D. Levine

The aim of the multisensor perception project is to investigate how models of the environment can be extracted from sensor data. Because one must first understand how sensors can be used to perceive the world, the project is a necessary step in building robotic systems that can interact with and adapt to the environment. This perception is concerned not only with identifying objects, but also with modeling their physical properties. For example, to grasp an object or to avoid colliding with it, a robot must know something about the size and shape of the object. But sensors such as television cameras and laser range finders can only measure such a size and shape from the information encoded in the sensor data. In addition, this perception must be robust and relate to a wide variety of physical properties. For this reason a multi-sensor approach has been adopted.

A5: 3-D Computer Vision: Sensing, Processing, and Interpreting Architectures

Project Leader: Dr. D. Poussart

The research program is focused on 3-D computer vision, a key component in the development of advanced manufacturing capabilities. It deals with three areas: (1) the study of 3-D sensing and the development of intelligent sensors, (2) the processing of 3-D images with emphasis on 3-D view integration, 3-D texture analysis, and the fusion of range and other visual information, and (3) the application of computer vision in industry and biomedicine.

The main objective of the program is the development of robust vision components and systems which can be implemented and used in industry: it hinges, in part, on applying principles of organization known in biological vision.

A6: Active Vision in Mobile Robot Navigation

Project Leader: Dr. J. Tsotsos

The goal of this project is to develop the active vision technology required for mobile robot navigation and exploration of unknown environments. According to the active vision paradigm, the process of acquiring and interpreting images is a closed loop process, in which information extracted from the images is used to control image acquisition, for example by controlling the parameters of the camera, and early processing. We will use mobile robot navigation as a test-bed for developing new active vision systems. At the same time, we will develop spatial representations and reasoning mechanisms for mobile robot path planning and the construction of cognitive maps through exploration. As part of the active vision research, we will develop algorithms for the robust fusion of information obtained from images through different types of processing, such as depth from stereo, motion, or focus.

The proposed IRIS project is complemented by a companion PRECARN project, entitled "Mobile Robotics," to be conducted jointly with Ontario Hydro and the Atomic Energy of Canada Ltd. The emphasis of that project will be on the implementation of a working prototype for operation in hazardous environments. On the contrary, the proposed IRIS project will emphasize basic computer vision research and will use mobile robot navigation as a test-bed for new results in active vision.

A7: Generation and Auditory Perception of Speech Signals

Project Leader: Dr. R. De Mori

The purpose of the proposed research is to investigate new theories for modeling speech units, lexical access in continuous speech dialogue between humans and robots, and supporting architectures. General paradigms for information access via voice will be investigated with particular attention to ways to extract relevant information from utterances which have free syntax form and contain words beyond the standard vocabulary. A mathematical mode of the ear coupled with multilayered networks will be developed for such a purpose.

Improving intonation in speech synthesis is another objective of the research. Ways to extract linguistic information from a text to be spoken which will lead to improved intonation models will be investigated.

Area B: Knowledge-based Systems

B1: Design of Large Information Systems

Project Leader: Dr. J. Mylopoulos

The aim of this project is to develop a suitable computer-based environment for the conception, design, implementation, and maintenance of large information systems, etc. There are four major research groups participating in the project, each focusing on a different aspect of information system design, namely rapid prototyping, software reusability, metasystem facilities, and integrated man–machine information systems. The project provides for the development of prototype environments which offer the features under consideration and also for the study of how these features facilitate the development of large information systems.

B2: Using Connectionist Learning for Adaptive Interfaces

Project Leader: Dr. G.E. Hinton

Simulated neural networks are composed of simple neuron-like computational elements that are richly interconnected. Learning procedures modify the connection strengths and thus allow the networks to be trained to interpret incoming data and to generate complex responses. In the last few years, there have been major advances in the area of learning techniques that allow the networks to develop their own internal representations, but these new learning procedures have not yet been applied to an important class of applications. When a person must provide high bandwidth control to a complex physical device, it is very important to use a "compatible" mapping between the person's movements and the behavior of the device. With most devices the mapping is fixed by the designer. It may be possible, however, to build devices in which the mapping adapts automatically, so that a person can control the device in the way that he/she finds most natural. Such adaptive interfaces could be used both for designing a compatible mapping and for allowing this mapping to adapt to different individual users. The research will investigate the use of automatic learning procedures to design compatible mappings.

B3: Database Techniques for Knowledge-base Management

Project Leader: Dr. A.O. Mendelzon

Consider a sensor-base entity employing a knowledge base to direct the operation of some system in the face of changing circumstances. What are the requirements for this knowledge base? First, it needs to store large numbers of atomic facts about the domain and make the relevant facts available to its client entities efficiently in response to queries. In robotics applications, these facts will usually include large amounts of spatial and temporal information. In addition to atomic facts, the knowledge base needs to store more general information, such as rules and constraints, and to be able to draw plausible and timely inferences from the rules, the constraints, and the facts, even when the available knowledge is insufficient to fully warrant such inferences from a logical point of view. To keep up with the changing environment, the knowledge base must be kept up to date by processing and incorporating the facts reported by the entity's sensors, dealing with any inconsistencies between the new knowledge and the old. Since many tasks will require multiple cooperating entities, the knowledge base must be concurrently accessible and capable of arbitrating the conflicts inherent in concurrent consultation and update. Finally, the integrity of the knowledge base should be resilient to hardware or software failures to provide reliable operation.

The focus of this project is the application and extension of techniques and results from database technology to the construction of systems that can manage such knowledge bases. In particular, we will work on the following topics: knowledge representation languages, graph-based query languages, efficient query processing and specialized reasoning strategies, user interfaces, concurrency control, programming languages for knowledge bases, and geometric and temporal data management.

B4: Design and Human Interfaces

Project Leader: Dr. N. Cercone

The Design of Human Interfaces Project is intended to be part of a broader project area entitled Knowledge-Based Systems in a Network of Centres of Excellence designated as the Institute for Robotics and Intelligent Systems (IRIS). IRIS will also include networked project investigations into computational perception and intelligent robotic systems. The Design of Human Interfaces Project will concentrate upon the investigation and integration of modes of human interaction with information systems, including graphics, text, and speech/natural language. Information technology is sufficiently available and so diverse in appearance that a cohesive national effort to discover and use what is best in current research on human interfaces may provide Canada with a distinct and very useful topic of expertise. Design, defined as the interaction between understanding and creation, must be understood to provide the theoretical basis that guides technological development of human interfaces. Principles of human interaction exist which have not yet been applied to make the abundance of information systems which already exist readily accessible to users. The well-designed and integrated interface will effectively exploit its knowledge of the domain and reasoning capabilities to facilitate interaction with users.

B5: Foundations for Reasoning Systems

Project Leader: Dr. R. Reiter

Canada has a significant research expertise in the foundations of automated reasoning, including world-class expertise in deduction and reasoning about time and action, default reasoning, analogical reasoning, explanation-based learning, hypothetical reasoning, and probabilistic reasoning.

The goal of the project is to actively promote and encourage the exchange of ideas and experience to provide a deeper understanding of a possible unified theory of automated reasoning which attempts to integrate reasoning theories for action, time, defaults, beliefs, analogy, learning, hypotheses, and probability.

The ultimate goal is to reduce the gap between these different theoretical models of aspects of common sense reasoning and to encourage work toward an integrated theory which will point the way to a general theory of artificial intelligence representation and reasoning.

B6: Engineering Applications of Constraint-logic Programming

Project Leader: Dr. M.H. vanEmden

An effective way to make computers more "intelligent" is to make it unnecessary to specify how a task is to be completed, so that one only has to say what needs to be done.

A method for achieving part of this objective is to program in terms of "constraints," where one only specifies properties the desired result of the computation should have. Logic programming, the method on which the Prolog programming language is based, is basically in terms of constraints. However, Prolog only covers symbolic constraints, to be satisfied by search.

A new development, constraint-logic programming, removes this limitation by allowing logic programs to include numerical constraints to be solved by special-purpose methods. The proposed project is to investigate applications in engineering, such as computer-aided design and the control of robot arms.

Area C: Intelligent Robotic Systems

C1: Telerobotics: Dynamics and Control

Project Leader: Dr. Y. Stepanenko

The main goal of this project is to develop intelligent control systems which provide effective and accurate telerobotic dynamic interaction with unknown and unstructured environments. This requires the development of new control methods for optimal, adaptive, and force control of compliant and flexible robot arms, and development of a suitable computer architecture for intelligent robot control. The project takes an integrated approach to these problems focusing on the dynamic aspects of intelligent robot behavior. An important point of the proposed project is the development of the Telerobotic Experimental Station (TEST), which will provide the experimental base for implementation and testing of new ideas, methods, and design solutions and will serve as a link for technology transfer. The applied results of this project will contribute to technological progress in such strategic areas as telerobotics (ocean rigs and northern development, forestry, mining, and space programs) and health care (prosthetics, exoskeletal structures, etc.).

C2: Simulation, Control, and Planning in Robotics

Project Leader: P.R. Belanger

This project is concerned with the operation of robotic systems. That operation is organized as a hierarchy. The lowest levels address questions of robot control of both force and position. This reburies mathematical modeling and simulation, in addition to measurement and the design of feedback control strategies. The upper levels are concerned with the planning of robot motion (or trajectory) and of tasks, through decomposition into subtasks. This project requires the construction of experimental test beds and the development of sophisticated computer systems.

C3: High-performance Robots: Design and Architecture

Project Leader: I.W. Hunter

The purpose of this project is to design, build, and analyze, theoretically and experimentally, high-performance robots. Such robots require fast, accurate, light-weight, direct-drive motors, novel optimized limb designs, and computer control architectures having floating-point speeds hundreds of times greater than those used in existing commercial robots. Our approach is not to build robots via the usual "shopping catalogue" or "off-the-shelf" method as this inevitably results in low-performance designs that are impossible to control adequately and that lack both accuracy and speed. We are consequently pushing the limits of current technology and where necessary developing new technologies.

C4: Intelligent Robot Planning and Execution Systems

Project Leader: H.A. ElMaraghy

This project aims at researching various aspect of action-plan generation and sensory-based execution and monitoring for intelligent robots. It consists of three interconnected modules: (1) Task Level Robot Planning, (2) Robot Path Planning, and (3) Multisensory Multirobot Control.

Intelligent robots operating in unstructured environments must be able to plan their actions and replan, in real time, to adapt to changing goals and conditions. Such ability requires reasoning and sensory-based control which take into consideration the uncertainty and incomplete information present in unstructured domains (e.g., underwater, outer space, and forests) and loosely structured domains, e.g., unfixtured flexible assembly). Professor H.A. ElMaraghy's objective at McMaster University is to provide a new environment for automated robot planning including an object-oriented geometric database and expert task and motion planner linked to a knowledge base and a run-time monitoring system. In addition to planning the actions of the whole robot arm, automatic grasp planning is also important. The generation of grasp configurations, i.e., points of contact with objects and corresponding grasping forces, and grasping manipulation with two or more fingers will be investigated.

Execution of generated plans of robot actions often requires careful robot trajectory planning to avoid collision with obstacles. The objective is to develop strategies for the automatic determination of feasible, optimum, and collision-free robot trajectories. A realtime trajectory optimization strategy to continuously update planned robot path, based on the current position and constraints in an unstructured environment, would also be researched by Dr. R.O. Buchal of The University of Western Ontario, London, Ontario. The robot system dynamics will be fully incorporated into the path planning problem formulation to take into consideration actuators torque and elapsed time. Research has already been started in this area by Professor W.H. ElMaraghy at The University of Western Ontario, to address the computer simulation of the dynamic behavior of multibody systems and development of a formalism for deriving the governing equations from system topology, mechanical properties of the bodies, and body interconnections.

Future robots working in an unstructured environment will also rely heavily on the efficient integration of sensory signals and control schemes. Professor M.M. Bayoumi at Queen's University, Kingston, Ontario, proposes to develop a robot control test station which is portable, modular, easy to modify, and easy to use. This would allow researching sensor integration and different levels of the control hierarchy. The overall modules in the station would include the servo level, the trajectory generation level, as well as low-level task planning. The sensors are classified into a hierarchy: low level such as position and force sensing, medium level such as range and tactile sensing, and high-level sensing such as vision. Such a test station should allow research to be carried out on hierarchial and intelligent control of robots with feedback from a bank of sensors.

C5: Mechanics, Control, and Design of Dexterous Robots/Environment Interfaces

Project Leader: A.A. Goldenberg

Robotic-based automation of processes in the industry and the service sector presents a formidable technical challenge to the engineering and scientific community. A key element of the advance of this area is the development of versatile interfaces which will extend the present capability, and thus improving the usefulness of robot manipulators for the execution of tasks in various environments. Furthermore, improvements in the technology of robotic manipulators are necessary to further reduce the costs of automation thereby maintaining a competitive edge within the Canadian engineering sector.

A multidisciplinary team of researchers of international stature in the Faculty of Applied Science and Engineering at the University of Toronto, Toronto, Ontario, will lead this research effort into both theoretical and applied aspects of the mechanics, design, and control of versatile dexterous robot interfaces. Drawing on a broad base of technical expertise, the team will examine technical areas that currently present barriers to the advancement of the state of the art in the science of robotics. Fundamental and applied research in key areas will permit the development of truly versatile robot interfaces with great potential for economic benefits.

C6: Teleoperation: Measurement of Machine Parameters and Human/Machine Performance

Project Leader: P.D. Lawrence

Heavy equipment controlled by human operations is used in important Canadian industries such as mining, forestry, agriculture, and construction. Examples of this equipment are excavators, cranes, log loaders, combine harvesters, and feller-bunchers. Virtually none of this equipment benefits from advanced computer controls of the type used in the manufacturing industries.

On-board computers can be used to assist the operator in carrying out the control of the machine by improving accuracy, repeatability, reliability, and safety and reducing fatigue. Preliminary studies have shown that novice operators can complete tasks more rapidly and more accurately with the computer-assisted controls, and that some experienced operators are also able to improve their performance.

This project concentrates on three important problems in the development of the technology. The first is the development of means to measure and transfer the characteristics of a particular machine to its control computer. The second is the development of important sensors that can measure operation and production variables for the computer. The third is the measurement of performance of the machine and the human.

The main industrial linkages of this project are in the forest and agricultural industries although interest has been expressed in the mining and construction industries.

C7: Dynamics and Control of Complex Space Manipulator Systems

Project Leader: P.C. Hughes

The use of remote manipulator systems for space applications exemplifies the ever-increasing complexities of robotic systems in general. The robotic systems of the future will be topologically elaborate and structural flexibility will no longer be an ignorable issue. The problems they will pose for dynamics modeling and simulation, path planning, and control will be formidable. It is unlikely that present-day techniques alone will be adequate to meet the challenges that lie ahead. New approaches must be explored.

This project is accordingly dedicated to the study and development of dynamics analysis methods and control techniques for complex robotic systems for space applications. Although the primary motivation derives from the use of robotics in space, it must be emphasized that the fruits of this project will also benefit the greater population of robots in industry.

The development of dynamics models will concentrate on complex topology, structural flexibility, and computational efficiency. As well, the architecture of revolutionary "parallel-processing minisupercomputer" will be exploited in designing advanced simulation algorithms for general multibody systems. In the area of control, the use of neural networks, a concept from the field of artificial intelligence, will be investigated. As their name may imply, neural networks attempt to replicate the computational scheme of the human brain. Although neural network theory is relatively new, the advances being made will surely have a major impact on the robotics community. Finally, to render the project as a coherent and complete package, the theoretical work will be validated in the laboratory using Radius, the UTIAS Space Robotics Laboratory Facility.

C8: Dynamic Design and Control of a Large Class of Space and Ground-based Flexible Manipulators

Project Leader: Dr. V.J. Modi

The ground-based robot systems tend to be relatively small in size, simple in design, and essentially rigid. Furthermore, they operate in the one-g continuum atmospheric environment where structural behavior and its control are relatively well understood. However, with the space-based flexible robot systems operating in the free molecular, microgravity, magnetic, and solar radiation environment, the problem acquires the challenge of a higher order of magnitude.

The challenge is further accentuated by the fact that the space environmental effects cannot be simulated even in the most sophisticated earth-based facilities. Thus, one is faced with the situation where theoretical procedures will be increasingly relied upon in the design of the space-based robots. Furthermore, we will have to find answers to these complex problems as Canada is participating in the U.S. sponsored Space Station program with the prime responsibility of design, development, and operation of a flexible robot manipulator referred to as the Mobile Servicing System (MSS). Dynamics and control of flexible robot manipulator, mobile and operating on the flexible station, carrying flexible payload, demand modeling procedures not encountered so far.

The comprehensive proposal aims at laying a sound foundation for the design of such next-generation robot manipulators. It attempts to achieve this through key information concerning interaction between librational dynamics, flexibility, MSS maneuvers, and control strategies which are fundamental to all the phases of the Space Station: construction, operation, servicing, and utilization.

C9: Dexterity in Teleoperation

Project Leader: Dr. J. Hollerbach

This project has as a goal to capture and utilize the dexterity of the human hand in a teleoperated system. Two teleoperated systems are proposed: (1) a force-reflecting master glove to control a multifingered robot hand in grasping and manipulating arbitrary objects, and (2) a fine-motion telerobot for high-performance tool use. Human factors studies are included as an essential component to provide a scientific basis for evaluating the performance of the master-slave units and to set design goals.

For more information, please contact:

Dr. A.E. Collin
Director, IRIS
PRECARN Associates Inc.
30 Colonnade Road, Suite 300
Nepean, Ontario
K2E 7P4
Telephone: (613) 727-9576
Fax: (613) 727-5672

Government Support Programs for AI R&D in Canadian Industry

A summary of the assistance programs of the Federal Government to encourage technological development in Canadian industry is shown in the accompanying table (Table 4). Incentives available to universities and nonincorporated businesses are excluded from the table. Also excluded is federal spending for targeted basic and applied research, a substantial portion of which is contracted out to the private sector, universities, and others.

Detailed program descriptions are available from the contacts identified in the table. Alternatively, prospective applicants can contact Industry, Science and Technology Canada (ISTC) regional offices or the field advisors of the National Research Council (NRC) located in major Canadian cities.

A brief description of the federal programs follows. With the exception of the AI R&D Fund, these programs are not exclusively directed at AI. However, they can be used to support industry performed AI R&D.

Artificial Intelligence Research and Development Fund (ISTC)

The AI R&D Fund exists to encourage federal government departments and agencies to develop and use AI systems in their operations. The Fund will provide up to matching contributions and is a $10 million five-year program. *Only federal government departments and agencies can apply to the Fund.*

Projects must have a significant level of contracting out to qualify for support. As well, projects must offer the potential of productivity gains within the sponsoring department and have potential applications in the private sector that will help improve the competitiveness of Canadian industry.

The Department of Supply and Services manages the process of contracting out. Private sector firms interested in bidding on contracts resulting from this program should contact DSS Science Branch. Each sponsoring department is responsible for the management of its own project.

An initial allocation from the Fund of approximately $2 million over four years has been made; this will result in a total of over $4.6 million being devoted to AI technology development and application as a result of matching contributions. A detailed list of projects approved to date is included in this report.

Strategic Technologies Program

This program provides financial assistance for the development and application of advanced industrial materials, biotechnology, and information technologies. It covers up to half the eligible costs of projects in leading-edge technologies undertaken by research and development and technology application alliances.

Two types of alliances are supported under this program:

- *Research and Development Alliances*: a company sharing risks and working with one or more companies, universities, or research institutes on precommercial R&D to develop the technology base needed for a range of new or improved products and processes.
- *Technology Application Alliances*: at least two companies engaged in complementary activities and one or more other organizations sharing risks and working jointly on precommercial technology development and related studies with the following objectives: to determine the production, economic, or market feasibility of new technological products or processes; to develop standards needed to permit technology applications; or to demonstrate to potential users in Canada the feasibility of leading-edge technology. Partners in this type of alliance can be users or developers of technology.

Microelectronics and Systems Development Program (ISTC)

The Microelectronics and Systems Development Program (MSDP) covers up to 50% of eligible costs incurred in developing innovative microelectronic and information technology systems that will enhance the international competitiveness of companies in manufacturing, processing, or service industries. It supports R&D concerned with innovative microelectronic components and underlying design and processing technologies for use in advanced systems or subsystems. Assistance is also provided for systems development R&D that

Table 4. Federal government support programs for AI R&D.

Program	Objective	Budget	Department	Applicants	Contact and Telephone
AI R&D Fund	To finance AI application projects in the Federal Government. Costs are shared with participating departments. The fund expressly support projects which are contracted out to the private sector or universities.	$10 million (5 years)	ISTC	Government departments	E. Raynes (613) 993-0303
Strategic Technology Program (STP)	To develop industry alliances in cooperative R&D and strategic technology applications. Industry-led alliances can involve universities and research institutes. R&D projects must be pre-competitive and application projects pre-commercial.	$30 million targeted to information technology (5 years)	ISTC	Alliances, consortiums, joint ventures and other collaborative ventures in industry.	André St. Pierre (613) 954-3470
Microelectronics and Systems Development Program (MSDP)	To assist in the development of advanced microelectronics and information technology systems. (Note: MSDP does not support standalone software projects. To qualify, AI activities must be an integral part of a technology or system development project.)	$60 million (5 years)	ISTC	Incorporated businesses	André St. Pierre (613) 954-3470
Technology Outreach Program	To support technology centers in technology diffusion initiatives	$15 million/ year	ISTC	Technology Centres	Elizabeth Payne (613) 954-3449
Defence Industry Productivity Program (DIPP)	To assist defence-related industries in technology development and production. The program provides innovation assistance for defence R&D. AI projects qualify for support provided, they satisfy defence-related needs and applications.	—	ISTC	Defence-related enterprises	P. Wilburn (613) 954-3338
Industrial Research Assistance Program (IRAP)	IRAP services include cost-sharing programs to support innovation activities for new products and processes.	—	NRC	Incorporated businesses	IRAP office (613) 993-2012

is based on advanced microelectronic and information technologies with applications in manufacturing, processing, or service industries. The development of standalone software packages generally does not qualify under MSDP. For AI projects to qualify for support, the software development activities must be an integral part of the development of a technology or system qualifying for support under MSDP.

Technology Outreach Program (ISTC)

The Technology Outreach Program (TOP) improves the productivity and competitiveness of Canadian industry by supporting technology centers whose national activities and services accelerate the acquisition, development, and diffusion of technology, especially in small- and medium-sized enterprises. The program provides both start-up and, in some cases, sustaining support to nonprofit technology centers. Up to 50% of the forecast operating costs of a proposed center, averaged over a five-year period, is available.

Industrial Research Assistance Program (NRC)

The Industrial Research Assistance Program (IRAP) increases the calibre and range of industrial research and development in Canada while encouraging the use of available technology. The program provides technical and financial support to firms for approved industrial research projects of high technical merit with good potential for commercial success. Such projects should be aimed at innovative products or processes that are realistic to the company involved and that provide significant economic or social benefits for Canada. Projects over $100 000 must involve collaboration with and/or technology transfer from a government, university, or foreign laboratory.

Regional Programs

The Federal Government also administers a number of regional programs which may complement the research and development activities of Canadian AI firms. These include:

- *Western Diversification Program*: This program helps businesses that are diversifying the Western economy by developing new products, new markets, new technology, import replacement, or industry-wide productivity improvements. The department shares the risks of these new ventures, usually by providing repayable interest-free contributions. Applicants must have significant equity in each project. If a project could proceed without WDP support, it will not be eligible. The department also plays a major pathfinding role, putting Western businesses in contact with other sources of financing.
- *Atlantic Canada Opportunities Agency*: This agency is directly responsible for federal policies and programs pertaining to small- and medium-sized businesses and industrial development in Atlantic Canada. It offers financial assistance for a variety of activities including feasibility studies, market research, and technology transfer. Up to 75% of the cost of feasibility studies and technology development and transfer activities may be available. As well, up to 60% of the cost of researching and developing a new product, process, or service may be provided.

Projects Co-sponsored by the AI R&D Fund

Development of a Medical Device Risk Assessment Decision Support Expert System (MEDRISK)

Sponsor: Health and Welfare Canada, Bureau of Radiation and Medical Devices (BRMD)

Contact: Fred Fiksel (613) 954-6703

The MEDRISK system will provide BRMD with an expert-system-based tool to perform comprehensive risk assessments on generic classes of medical devices and assist in identifying suitable protocols for the regulatory review of device-related information. The BRMD maintain responsibility for regulating the sale of medical devices in Canada, so as to ensure their safety and efficacy. The large and rapidly growing number of devices on the Canadian market, together with the increasing level of device sophistication and rapid rate of technological change, strains the limited resources which BRMD devotes to the regulatory review process. It is expected that use of the MEDRISK system will facilitate more efficient and effective resource utilization, as MEDRISK will assist BRMD in conducting reviews which are more specifically tailored to the key risk-related characteristics of devices. Such tailoring will enable BRMD to shorten the regulatory review process for specific devices and broaden its scope of review to include a greater number and variety of devices. The foregoing should provide Canadians with quicker access to innovative and life-saving technologies, help industry in dealing with rapid technological obsolescence, and better ensure device safety and efficacy. The MEDRISK system is being developed with the assistance of the U.S. Food and Drug Administration. Contributions from the Fund to this two-year project will match that of Health and Welfare Canada.

Expert Systems for Terminal Weather Forecasts — Phase II

Sponsors: Transport Canada, Transportation Development Centre; Environment Canada, Atmospheric Environment Service; Department of National Defence

Contact: Howard Posluns (Transport Canada) (514) 283-0034

This proposal is for an advisory expert system to assist terminal weather forecasters in providing accurate low-cloud forecasts for civilian and military aviation use. This deep knowledge expert system will simulate qualitatively the physics of certain phenomena which the forecaster uses implicitly for his predictions. Such a system could improve the quality and reliability of terminal forecasts available, allowing aircraft to operate more safely and efficiently. The artificial intelligence technology to be developed in this project would be in high demand worldwide for similar applications. Weather forecasts for aviation are very important for the safe and efficient operation of civilian and military aircraft. These forecasts are the basis for planning safe and efficient flight plans which must include contingencies for local weather conditions. Terminal weather forecasters are highly skilled and must perform a considerable volume of work each day. Contributions from the fund will match the contributions of the three sponsoring agencies. Transport Canada is taking the lead for the project and is making the largest contribution.

RAISON/AI: An Expert System for Regional Assessment of Water Resources

Sponsor: Environment Canada, National Water Research Institute

Contact: David Lam (416) 336-4916

The RAISON/AI project will attempt to expand a successful existing regional assessment computer system by utilizing artificial intelligence technology to deal with the fuzzy data and heuristic reasoning required to make accurate regional assessments. Environment Canada has demonstrated success with a computer-based system to perform some of the data management required to integrate the information on a regional basis. The RAISON/AI system has the potential to improve greatly Canada's capability to assess environmental factors on a regional basis. The assessment of the impacts of various environmental factors is difficult even on a local basis. Determining those impacts on a regional basis is extremely difficult, especially when much of the required information is qualitative rather than quantitative. Prototypes of the system have received considerable interest from foreign countries, which could represent a significant export market for the technology. Environment Canada's contribution to this four-year project will be matched with contributions from the Fund.

Development of a Microcomputer-based Expert System for Gold Mill Effluent Treatment Plant Design

Sponsor: Environment Canada, Wastewater Technology Centre

Contact: John Neate (416) 336-4599

This expert system will integrate the knowledge required to make assessments of complex gold mill effluent plant operations into a user friendly expert advisor which could be used both by regulators (to aid in the approvals process) and the mining companies (to design efficient, compliant plants). Gold mining has considerable importance to Canada. Our gold reserves continue to generate significant economic returns. The gold milling process, however, involves the generation of considerable quantities of potentially hazardous effluent. The design of effective and efficient effluent treatment plants not only will ensure that impact on the environment will be minimized, but will also result in savings to firms through reduced down time and reduced costs for compliance with regulatory requirements. The design of such facilities requires a wide range of information, including process factors and local knowledge of the gold deposits and expected effluent streams. Environment Canada must be able to assess potential designs to ensure regulatory compliance, while industry would like to choose a plant design which will ensure compliance at the lowest cost. Fund contributions will match the contribution of Environment Canada to this project. Significant funding is also being contributed by the Ontario government. The Canadian gold mining industry is expected to contribute.

System of Hierarchical Experts for Resource Inventories (SHERI Project)

Sponsor: Energy, Mines and Resources Canada, Canada Centre for Remote Sensing

Contact: Dr. David Goodenough (613) 952-2760

The Canada Centre for Remote Sensing will be developing an artificial intelligence system which would integrate data from remote sensing with a variety of other geographic information systems to perform forestry inventories. This approach would be a considerable improvement over the existing situation, which suffers from a shortage of the highly trained experts needed to interpret the data. Canada is such a large country that monitoring the status of our resource supplies is an enormous and expensive task. Management of Canada's natural resources is a priority, especially when supplies are dwindling and there are concerns for the state of our environment. Interest in expert systems for forest inventorying is high. The British Columbia Ministry of Forests is cosponsoring this project. This project will be undertaken by a government–industry–university team which will pool their expertise to develop a unique and valuable system. Contributions from the Fund will cover approximately 37% of the total cost of this 3½-year project.

Feasibility Studies Co-sponsored by the AI R&D Fund

In addition to the five systems development projects approved, feasibility studies for three projects were approved. These studies are aimed at demonstrating the applicability of AI technology to particular problems. The three studies are:

Machine-aided Translation and Bilingual Synthesis of Natural Language Texts

Sponsor: Department of Communications,Canadian Workplace Automation Research Centre

Contact: Pierre Isabelle (514) 682-3400

Intelligent Training Environments for Case Studies in Public Administration and Expert System for Case Writing

Sponsors: Department of Communications, Canadian Workplace Automation Research Centre; Canadian Centre for Management Development

Contact: Jocelyn Picot (514) 682-3400

SEPA —Strategic Environmental and Process Assessment System

Sponsor: Environment Canada, Conservation and Protection

Contact: Abe Finkelstein (819) 953-1110

For more information, please contact:

André Dubois
Deputy Director, Technology Management
Information Technologies Industry Branch
Industry, Science and Technology Canada
235 Queen Street
Ottawa, Ontario
Telephone: (613) 954-3452

Natural Sciences and Engineering Research Council

The Natural Sciences and Engineering Research Council (NSERC) is a federal granting agency with the mandate to "promote and assist research in the natural sciences and engineering other than the health sciences." Council focuses the bulk of its activities on the support of university-based R&D and research training and is the largest single source of funds in Canada for this purpose. A more recent thrust has been to forge closer links between the university research community and other sectors to promote collaborative R&D activities and technology transfer.

For the current fiscal year (1990–91), NSERC's total grants and scholarship budget will exceed $400 million. It provides research or training support to some 17 000 individuals.

The sub-objectives of Council's program of Grants and Scholarships in Aid of Research are:

- to support a diversified base of high quality research in the natural sciences and engineering

- to assist in the provision and development of highly qualified personnel

- to promote and support targeted research in selected fields of national importance

- to forge closer links between the university research community and other sectors of the economy

To achieve these sub-objectives, NSERC's programs are divided into three major activities:

- Research Base Programs (including Operating Grants)
- Scholarships and International Programs
- Targeted Research (Strategic Grants and Research Partnerships Program)

The NSERC responds to requests for funds from a broad cross section of university researchers. Applications are reviewed by committees of peers, and awards are made according to the criteria of the particular awards program and within the budget allocated to the program from the total funds available.

For certain specific research areas, NSERC's ability to report its expenditures is limited by our coding system and by grantees' application of this system to their research. Therefore the figures available are for selected programs (Research Base and Targeted Research) and only provide an indication (and likely an underestimate) of NSERC's total support of artificial intelligence. With these caveats in mind, in 1989–90 NSERC expended in the order of $3.6 million on artificial intelligence research.

The Research Base Programs, and especially its major element the Operating Grants program, provide support for a diversified base of excellent research across the country. The Research Base Programs are aimed at providing continuity of support and at creating the core environment for the training of graduate students. Activities supported by other NSERC programs or by other organizations at Canadian universities, in areas such as AI research, are built on the research capabilities fostered by these programs.

The Targeted Research Programs consist of Strategic Grants and Research Partnerships. Strategic Grants provide support for the initiation or the acceleration of substantial projects of finite duration in certain areas of national importance. The levels of funding provided in the Strategic Grants Program permit researchers in artificial intelligence to make significant contributions in an area where the cost of research can be high. The program encourages multidisciplinary group research and attracts researchers from a broad range of disciplines. Support for research in artificial intelligence is provided for under several of the strategic areas, particularly information systems and manufacturing systems.

The Research Partnerships Program consists of various R&D activities jointly financed by NSERC, industry, and other federal government departments. They include cooperative R&D grants for research projects or programs carried out by university researchers in cooperation with industry and/or the federal government; industrial research chairs which help universities to achieve the critical mass of expertise essential for major scientific or technological research endeavours in collaboration with industry; and various activities aimed at improving technology diffusion through the exchange of research personnel between universities and industry, and through the support of workshops, seminars, and university–industry affiliations. Recently supported initiatives include the creation of an Industrial Research Chair in Robotics at McGill University.

With the Medical Research Council (MRC) and the Social Sciences and Humanities Research Council (SSHRC), NSERC administered the recent competition for the national Networks of Centres of Excellence (NCE). The Institute for Robotics and Intelligent Systems (IRIS), which brings together 18 universities and PRECARN Associates Inc., was one of the 15 networks selected for funding. Under the NCE program, NSERC will grant $23.8 million to IRIS over the next four years.

Governments of all industrialized countries are quickly becoming aware of an increasing demand for more highly qualified scientific and engineering personnel. In R&D-intensive sectors, such as AI, an adequate supply of research personnel is critical to the Canadian economy, which is becoming more knowledge intensive. Therefore, in its recently published strategic plan, NSERC has placed a major emphasis on its role in meeting the demand for highly qualified personnel. The NSERC supports research training directly through its Scholarships and International Programs, and also indirectly through stipends paid from Operating, Strategic and Research Partnerships grants.

For further information, please contact:

Natural Sciences and Engineering Research Council of Canada (NSERC)
Centennial Towers
200 Kent Street
Ottawa, Ontario
K1A 1H5
Telephone: (613) 995-6295

Section 4

AI in Universities, Government Departments, and Agencies

Canadian Space Agency

Keywords: AI, planning, diagnosis, vision, tutoring, machine learning, MSS, SPDM

The Canadian Space Agency (CSA) supports a number of AI activities in the space sciences and applications, and through industrially based capital space projects. The majority of the AI Research and Development activity within the CSA is sponsored by the Space Station Mobile Servicing System (MSS) and the Space Station Strategic Technologies for Enhanced Automation and Robotics (STEAR) programs. In addition, the Directorate of Space Mechanics (DSM) has active in-house projects utilizing expert systems and neural networks.

The CSA-sponsored AI activity is briefly outlined with respect to each technology area. The MSS baseline activities are discussed in the first part, MSS evolutionary activities in the second part, and DSM in-house activities in the final part.

MSS Baseline Activities

The Canadian contribution to the International Space Station "Freedom" is the Mobile Servicing System. The MSS is a collection of robotic elements configured to support the assembly, maintenance, and servicing of the space station. The major components of the MSS are:

- Space Station Remote Manipulator System (SSRMS): a large robotic manipulator arm with seven degrees of freedom used primarily for handling large objects on station including the capture and docking of the shuttle orbiter.
- Special Purpose Dextrous Manipulator (SPDM): a smaller two-arm robotic device capable of performing robotic servicing tasks requiring dexterity.
- Mobile Servicing Centre (MSC): a mobile base for the MSS from which the robotic manipulators can be operated remotely and used to transport payloads around the station.

The basic or baseline capability of these components at the Initial Operational Configuration (IOC) of the Station will incorporate strategic technologies, including expert systems utilizing AI techniques. The majority of the baseline activity is contracted to the Space Station Industrial Team which includes Spar Aerospace, CAE Electronics Ltd., Canadian Astronautics Ltd., SED Systems Inc., MacDonald Detwiller and Associates Ltd., and IMP Group Limited.

The baseline MSS/SPDM systems have been designed with the objective of evolving toward autonomous operations. A command and programming language has been developed for MSS to implement automated procedures. This language uses AI techniques to accommodate the dynamic nature of the MSS environment and changes in the MSS. A collision avoidance system has also been developed which will provide MSS with an environmental model. An experimental machine vision system referred to as the "space vision system," which has the potential to provide closed loop control, has been developed and will be tested on a space shuttle mission in 1992 by a Canadian astronaut. These systems will comprise the core of the autonomous functions and are described below. Both the command and programming language and collision avoidance system have been developed to a prototype level and are under review for incorporation into the baseline MSS.

Procedural Reasoning System

This system will provide for capture of failure diagnosis knowledge in the format of procedures as opposed to rules. The system will execute in real time to aid MSS operators with determination of failure states of the MSS. Procedures are selected and executed in response to the sensed state of the MSS and established goals. A graphical procedure editor has also been developed for procedure generation. Shuttle RMS malfunction procedures have been implemented in the prototype system.

Planning System

A planning system is being developed initially to assist in the generation of plans for MSS operations. This system will be domain independent, address both high- and low-level plans, reason about resources and time, and will be implemented in a layered structure to permit future extensions. A knowledge acquisition tool (MEDUSA) has been created to assist MSS operations engineers with the analysis and formalized development of MSS tasks. The product of MEDUSA is a knowledge base of MSS operations, which is required by the planner but does not currently exist.

Trajectory Planning and Object Avoidance

The collision avoidance system developed by the Spar ATD group for the MSS baseline utilizes models of the Space Station and the MSS. Representation of the station in an octree format enables highly efficient checks for impending collisions to be performed. This representation also yields a free-space model which can be used for automated generation of collision free paths. Techniques for verification of the models have been developed. A prototype of this system has been developed which will be implemented on the SPDM ground testbed. Work is continuing on algorithms for realtime avoidance of collisions using the redundancy of the MSS manipulators.

Vision Systems

A fundamental component of the MSS Control Equipment (MCE) is an Artificial Vision Unit (AVU), which will play a major role in both manual and automatic control of the manipulators and cameras. The baseline specification states that the AVU will support the rendezvous and proximity operations including payload tracking, capture, and berthing. In one mode, the AVU shall provide data to the astronaut operator in the form of textual and graphical displays. In another mode, the AVU data shall allow the closed loop control of MSS manipulators for these operations.

Apart from the development of the experimental SVS for flight, the CSA has supported the design and assembly of a computer-based Machine Vision Development Platform (MVDP) that has been assembled by the Structures and Materials Laboratory of the National Aeronautical Establishment (now the Institute of Aerospace Research) of the National Research Council. In a joint research program, the CSA and the Division of Electrical Engineering (now the Institute for Information Technology), also of the NRC, have developed an intelligent model-driven variable resolution laser scanner. The CSA has also contributed to the development of a Stereovideo/Stereographic system sponsored by the Robotics Group of the Defence and Civil Institute of Environmental Medicine. It is anticipated that a coordinated program with contributions from each of these three groups and the Space Station Industrial Team will be formalized with the objective of developing prototype vision technologies that meet the baseline requirement of an autonomous sensor-based manipulator for MSS/SPDM. These vision technologies would support an AI system having realtime roles in automatic trajectory planning and execution, collision avoidance, and the health monitoring, fault detection/fault isolation system.

On-orbit Integration, Test, and Maintenance Information System (OITMIS)

This system has been developed by the University of Waterloo working in conjunction with the Spar ATD group to capture and provide electronic access to the MSS engineering database, in the form of documents, drawings, tests, operations, and maintenance procedures, etc. The system is intended to incorporate features to provide electronically the same functionality as a paper-based library (with a good library management system), equipped with photocopies and study areas.

Evolutionary Activities

With respect to an evolutionary capability of the MSS, the Canadian Space Agency is also active in a number of areas in AI research and development through sponsoring industry-based activities. The industry-based work takes place both in the Advanced Technology Development (ATD) Section of the Space Station Industrial Team and through the STEAR program.

The STEAR program was established to solicit the participation of Canadian companies outside the Space Station industrial team, as well as universities and research organizations, in developing strategically important automation and robotics technologies for the evolution of the MSS. The contracts are frequently carried out in two phases. The purpose of Phase 1 is to perform preliminary/feasibility studies, to prepare an exploitation plan and submit a detailed technical proposal for Phase 2. During Phase 2 the team must demonstrate proof of concept for a space and/or terrestrial application.

Expert Systems

The CSA has funded significant Expert Systems R&D activity in industry largely through the STEAR Program and the Space Station MSS Advanced Technology Development funding. A number of contracts are now active in various areas of expert systems development.

Diagnosis and Planning

Diagnosis and Planning have been identified as two of the areas which must have significant involvement of the human operator over the duration of a space system mission. At present there are five industry contracts in the diagnosis and planning domain. The five industry contracts are:

1. *Automated Expert System for Operations Planning*, Dynacon Enterprises Limited, Toronto, Ontario. A blackboard shell "AESOP" is used for the underlying structure in the planning system. Research includes user interface, human factors issues, resource management, and an explanation facility (Space Station STEAR funding $900 000 over 3 years).

2. *Automation of Operations Relating to Fault Detection/Diagnosis and System Health Monitoring and Failure Anticipation*, Centre for Frontier Engineering Research (CFER), Edmonton, Alberta. This system performs diagnosis and planning in a hierarchical manner using traditional realtime and a variety of rule-based and model-based AI techniques. It is intended to provide the operator with recommended corrective actions, for example, in the maintenance of semi-submersible drilling platform ballast control, in a variety of operational modes and weather conditions. It has the ability to operate in a very dynamic and time varying environment where degraded operation based on erroneous data is common (Space Station STEAR funding $900 000 over 3 years).

3. *Health Monitoring Power Management and Control System (HMPMCS)*, ISE Research Limited, Port Coquitlam, British Columbia. This system performs diagnosis and resource management (planning) as applied to a spacecraft electrical power system (space-based radar). The diagnostic system component combines realtime and AI techniques into a system capable of reacting to single faults in real time and multiple faults in an off-line manner. Realtime fault detection and isolation uses traditional techniques while the off-line system relies on both a rule-based and model-based approach at diagnosis. The resource management (planning) system component relies on a rule-based constraint resolution scheme (Space Station STEAR funding Phase 1 completed spring 1990, $190 000. Phase 2 $900 000 over 3 years pending approval).

4. *Automation of Operations Related to Planning under Stress, Automated Mission Planning Support and Problem Solving*, Softwords, Victoria, British Columbia. The contract has recently been modified from an expert-based training system to the development of a computer-assisted planning system. The system will test/evaluate candidate plans and then select the appropriate plan given the current system and environmental constraints (Space Station STEAR funding $900 000 over 3 years).

5. *Automation of Operations Relating to Automated Data Management Systems*, Spectrum Engineering Corporation Limited, Peterborough, Ontario. The system design relies on the use of Nial to manipulate and maintain the knowledge base and data structures to perform fault detection, diagnosis, failure prediction, fault history reporting, reliability analysis, and data presentation. Towed submersible cable faring maintenance is the example application being studied (Space Station STEAR funding $900 000 over 3 years).

Tutoring Systems

Automation of Operations Relating to Crew Training: Intelligent Tutoring Systems, Acquired Intelligence Inc, Victoria, British Columbia. Through the use of automated knowledge acquisition techniques and curriculum author, the expert/instructor develops a domain knowledge base and a tutorial course. The student then interacts through the tutor to develope their knowledge of the domain. This system includes an advanced user interface incorporating both hypermedia and user modeling techniques (Space Station STEAR funding $900 000 over 2.5 years).

Autonomous Robotics

Another STEAR Request for Proposal (RFP) deals with the development of technology to provide limited or supervised autonomous robotics capability on the MSS. Both teleoperated and autonomous robotics are baselined for the MSS SPDM. However, it is recognized that there will be a shift in emphasis from teleoperated robotics to autonomous robotics during the 30-year life of the Space Station Freedom. The field of AI shows promise to advance the technology towards autonomous robotics in its ability to model the environment; perform sensor fusion; plan actions/tasks, plan paths and avoid obstacles; support hierarchical architectures, communications, and scheduling; and manage faults.

A review of six detailed autonomous robotics technical proposals from the following firms is planned for the summer of 1990.

- Dynacon Enterprises Ltd., Downsview, Ontario
- Vadeko International Inc., Mississauga, Ontario
- Thomson CSF Systems Canada Ltd., Nepean, Ontario
- International Submarine Engineering Ltd., Port Coquitlam, British Columbia
- Robotics Systems International Ltd., Sidney, British Columbia
- MPB Technologies Inc., Dorval, Quebec

Phase 2 contracts will be subsequently awarded in the fall of 1990.

Trajectory Planning and Obstacle Avoidance

A STEAR RFP was closed at the end of June 1990. This RFP deals with the area of manipulator trajectory planning and obstacle avoidance under the requirement of realtime automatic control, with a minimum of human intervention. It is expected that advanced AI techniques will more efficiently handle the dynamic environment of the space station. The successful companies will explore the ability of these techniques to detect objects, identify possible dangerous conditions, plan manipulator trajectories, and respond, all in real time. Phase 1 contract awards are planned for the fall of 1990.

Vision Systems

In this STEAR initiative six Phase 1 contracts for enhanced vision systems (approximately $100K each) have been awarded to:

- Hymarc Ltd., Ottawa, Ontario
- Oerlikon Aerospace Ltd., Saint-Jean-sur-Richelieu, Quebec
- Optech Inc., Downsview, Ontario
- Sensor Adaptive Machines Inc., Windsor, Ontario
- Telepresence Research Ltd., Vancouver, British Columbia
- Virtek Ltd., Waterloo, Ontario

Each contractor will propose a vision system in competition for Phase 2 contracts that will be useful for object recognition/identification, 3-D surveying/surveillance, and realtime pose determination. Phase 2 contracts will be awarded to the successful companies at the end of 1990.

DSM In-house Activities

The Directorate of Space Mechanics (DSM) is pursuing a number of projects concerning the development of AI technology for space applications. The projects usually involve cooperative technical efforts between in-house, university, and industry participants.

Diagnostics and Planning

In the Computer and Intelligent Systems Group of the DSM, the thrust of AI research has been in the Expert Systems area as applied to planning and diagnostics. The internal research projects are the Case Based Planning System (CBPS) and the Multiple Fault Diagnostic System (MFDS).

- The *Case Based Planning System* uses case-based reasoning techniques to maintain a plan library and perform: plan selection, revision, execution, dynamic replanning, and plan evaluation. The system is presently applied to the spacecraft electrical power system domain and performs resource planning.
- The *Multiple Fault Diagnostic System* relies on model-based reasoning techniques and assumption-based truth maintenance to perform diagnoses in an environment containing multiple faults in an electrical circuit. In the presence of other faults, the MFDS detects faults, diagnoses the faulty component, maintains a fault history, maintains the model parameters, and recommends corrective action.

The Principal Investigator in this research is P.J. Adamovits.

Machine Learning

In the Dynamics Group of the DSM, the major thrust of AI research has been in the area of artificial neural networks as applied to nonlinear systems and manipulator dynamics. The prediction of nonlinear structural systems and chaotic time series has been modeled with neural networks using back-propagation learning. Learning performance has been increased by implementing an inherently parallel architecture on a multitransputer system. The inverse dynamics of a several degrees of freedom robot has been learned by a specialized neural network. This architecture can generalize to similar manipulator configurations, includes the basic properties of the dynamic equations, and permits cooperation between several networks at once. All of these factors contribute to the capabilities of the network and will allow transfer of the simulation software to a dual-arm realtime hardware testbed.

The Principal Investigator in this research is D.P.W. Graham.

Automated Validation of Thermal Mathematical Models (TMMs)

In the Applied Mechanics Group of the DSM, the thrust of AI research concerns the application of expert systems technology to the automated validation, in near-realtime during thermal testing, of the TMM for the hardware under test. The validation of TMMs constitutes a necessary but expensive part of a spacecraft program. Two significant components of the expense are (i) the time consumed by engineering personnel, following the test, in analyzing the data and in modifying the TMM to obtain corroboration with the data and (ii) the lost opportunity, after analyzing the data, to perform additional tests to investigate suspect data or unexpected behavior (the opportunity being lost because by this time the test article may have been removed from the chamber, if not already shipped elsewhere). The automation of the process to obtain near-realtime TMM validation will substantially reduce or eliminate these two components of expense.

The automated validation problem may be viewed as an inverse mathematical problem, in which the TMM temperatures (corresponding to test sensor locations) are known and the TMM coefficients (groups of design parameters) are to be determined. The knowledge for this problem is embodied in (a) the TMM, (b) the test data, and (c) a rule-based fuzzy logic reasoning structure which is designed to mimic the reasoning processes of thermal experts when they attempt manually to corroborate (a) and (b).

The Principal Investigator in this research is V.A. Wehrle.

Conclusions

The Artificial Intelligence Activities of the Canadian Space Agency have been briefly described. They include the technology areas of diagnosis, planning, tutoring, machine learning, vision systems, trajectory planning, and obstacle avoidance. The baseline configuration of the MSS/SPDM contains some AI and it is clear to the CSA that the future incorporation of advanced AI techniques developed through the evolutionary and in-house R&D activities will substantially improve the operation, planning and execution of both ground-based and on-orbit tasks.

Prepared by Dr. Steve MacLean, Peter Adamovits, and David Hunter.

Department of Communications, Canadian Workplace Automation Research Centre

Introduction

The Canadian Workplace Automation Research Centre, which opened on November 5, 1985, is part of the research program of the Department of Communications. CWARC acts as a catalyst for the growing use of telecommunications and computer technologies in response to the needs of knowledge workers in the office of the future. The majority of the research undertaken at the Centre is applied while a smaller proportion (20%) is dedicated to basic research areas. Artificial intelligence research and development projects are undertaken by the Advanced Technology Directorate, which is described in greater detail here. Other CWARC directorates include Organizational Research, Integrated Systems, and External Cooperation.

CWARC research is undertaken through joint ventures with partners from the private, public, and academic sectors to ensure the effective transfer and application of technology from university and government laboratories to address private sector business needs. This transfer is further assisted by the National Exchange Program, which is an innovative approach for the transfer of technology and the recruitment of scientific staff through personnel exchanges. Experts are seconded to CWARC for periods ranging from six months to possibly three years, to participate in research projects. The dynamic nature of research teams permits timely interventions in the rapidly changing field of new workplace technologies.

Artificial Intelligence Projects

There are two major groups working in the field of artificial intelligence in the Advanced Technology Directorate: the Computer-aided Translation Group and the Expert Systems Group. An overview of some of their major research projects is presented here.

Expert Systems

This area is comprised of three groups: expert systems for knowledge management, expert systems for know-how management, and expert systems for information management. Knowledge management research deals primarily with methods of knowledge acquisition and knowledge modeling. This section is headed by Joel Muzard. Know-how management expert systems are intended to serve as vehicles for education and training. Kim Dalkir is responsible for this area. Finally, information management expert systems can provide more advanced techniques for the automatic indexation of documents. This research is overseen by Jennifer Farkas.

The primary objective of these groups is to develop knowledge-based system applications for knowledge workers in offices. The field of office automation presents innumerable opportunities to implement artificial intelligence systems. We endeavour to implement a value-added approach whenever possible, namely, to add on an intelligent layer to existing office systems. In this way, the expert job aids are more likely to be used than if they were isolated, standalone systems. Above all, our philosophy is one of user pull as opposed to technology push — our research and development projects must address real user needs and provide actual business solutions at all times.

Some expert systems completed to date include QUARTZ, an intelligent assistant to help government analysts evaluate project proposals submitted for funding to the ERDA (Economic Regional Development Assistance) program. This system was originally developed for the Federal and Provincial Departments of Communication. It has subsequently been modified for use by the Transport Development Centre (QUARTZ II) and for internal use by CWARC (the PERIDOT system).

Another example is AMETHYST, an intelligent job aid, developed for pay and benefits clerks at the headquarters of the Department of Communications, used for Taken on Strength (TOS) transactions (hiring someone in the Public Service). The AMETHYST system consists of 13 expert systems, each corresponding to an area of pay and benefits (e.g., salary eligibility, medical insurances, income tax, etc.) as well as an additional knowledge base to drive the interaction with the user in an intelligent way. The system has been integrated with a database of data on employees and salary rate files. This database serves as a link to the MEMO system, which can scan information on documents (e.g., a ROST, Report of Staffing Transaction), place this data in the database, receive deductions and data from the expert systems, and then print out the appropriate and correctly completed forms.

These systems are operational prototypes that were developed on IBM PC (or compatible) computers using the ARGUMENT and DECIDEX development tools, with additional programming to further customize the application to meet user needs and to fully integrate them in user environments.

Some knowledge-base development projects currently underway include: ALADIN, an environment to assist knowledge engineers in knowledge acquisition and knowledge representation, and STAT, a tool to assist in the analysis of statistical data. A number of feasibility studies are also being carried out, including INDEX, a feasibility study on the use of artificial intelligence techniques to assist in the indexation of documents, and the INTELLIGENT COURSE SCHEDULER, a feasibility study to study the use of artificial intelligence for scheduling training courses in the Department of National Defence.

Finally, there are several exploratory research activities including the INTELLIGENT TRAINER project (with the Organizational Research Directorate) which explores the application of artificial intelligence to intelligent training systems for public service employees and the CONVERSATIONAL INTERFACE project, which seeks to determine the applicability of personal construct theory, viable system modeling, and conversation theory to interface design in intelligent tutoring systems.

Machine-aided Translation

If translation between natural languages were comparable to a simple transcoding process like that, say, between two formal languages, then the problem of automating human translation would have been solved years ago. But of course it is not. As every translator knows, to translate a text properly, it is necessary to *understand* it. Moreover, the level of understanding required for translation extends beyond the literal wording of the text itself, not just to what the author intended to say but may have formulated unclearly, but also his (and the translator's) extralinguistic knowledge as well as their ability to perform common sense reasoning. This is why machine translation is rightfully considered to be a branch of artificial intelligence and — if our goal is to fully emulate the capabilities of the human translator — an area of basic research.

The CWARC has a small team in machine-aided translation and an extremely ambitious program that includes both research and development components. The team is made up of four full-time researchers, two visiting scholars, and a number of part-time students; they bring together backgrounds in linguistics, computer science, translation, and artificial intelligence. This team is headed by Pierre Isabelle, a scholar with an international reputation in the field, who was formerly scientific director of the well-known TAUM group at the Université de Montréal.

The short-term objective of the group is to develop a specialized translator's workstation. The goal here is not to automate the actual translation process but simply to integrate a set of utilities that will help the human translator perform his tasks more effectively. Such utilities would include programs for bilingual word processing, the telecommunication of texts, building up personal and group glossaries, comparing two versions of a given text, accessing electronic dictionaries and term banks (either remote or local), in addition to various other writing aids. A modest first version of our workstation, based entirely on commercially available software, is currently being tested at the federal Translation Bureau. While we await the results of that trial, work has already begun on a more sophisticated second version.

In the mid-term, we hope to identify one or more promising sublanguages for which dedicated machine translation systems could be developed, on the model of the famous MÉTÉO system. We are convinced that given the current state of the art in machine translation, this sort of highly restricted application, designed particularly for text domains with simple syntax and large volumes, is the only way of simultaneously achieving high quality and full automation. A related type of technology that is also applicable to simple sublanguages is bilingual text generation. Rather than automating the translation of a previously written text, bilingual generation seeks to automate the writing process itself, as it is applied to two or more texts in parallel. This appears to be the ideal approach for stereotyped reports that synthesize trends in numerical databases. With support from the newly created AI R&D Fund, the CWARC is about to launch a project that will automate the drafting of several types of the bilingual summary reports produced at Statistics Canada.

Ultimately, however, the only way to resolve the worldwide crisis in translation is to develop machine translation systems that are more powerful, more "intelligent" than those available today. This in fact is the goal of the third component of our program: to apply the most recent artificial intelligence and natural language processing techniques to the particular problems of translation. To illustrate the progress we have made in our research, we have designed a prototype system, dubbed CRITTER, which translates agricultural market reports. Although CRITTER is an experimental system, implemented in Prolog on Sun workstations, it does possess extensive language skills. For example, the system is able to recognize and produce lexical and structural paraphrases. Moreover, CRITTER is the first machine translation system to be completely reversible: it can translate both from French to English and from English to French, using exactly the same grammars and dictionaries. If commercial machine translation systems ever hope to earn the respect of human translators, these are the sorts of properties they will have to incorporate.

For further information, please contact:

Marcel Drouin, Director
Advanced Technology Directorate
Canadian Workplace Automation Research Centre
1575 Chomedey Boulevard
Laval, Québec
H7V 2X2
Telephone: (514) 682-3400
Fax: (514) 686-1990

Department of Communications, Communications Research Centre

Speech Recognition Using Neural Networks

Research on automatic speech recognition at the Communications Research Centre Behavioural Research Division is aimed at speaker-independent recognition of continuous speech for interaction with new communications services. Recent progress at other laboratories has shown that sufficiently good performance may be obtained with a vocabulary of about 1000 words. The obtained high level of performance has resulted from improved speech representations and the availability of a large database of segmented speech suitable for training systems that learn. The work at the CRC uses the same database to represent the variability of speech in neural network models. Important research issues are how to modularize recognition systems for feasible learning of all speech sounds, how to integrate prosodic and phonemic information, and how to deal with the temporal aspects of speech.

The immediate objective is to develop a speech interface for new information services such as those provided by teletext or audiotext. Access to information would be facilitated for many people if spoken language could be understood by computer interfaces. For example, speech could be used to access the AIDS information base used to field test CRC's natural language query system (CHAT). Speech would also be a convenient means of interacting with terminal equipment in a mobile environment where manual control is difficult and unsafe. Broadcast audiotex services could provide current information of value to automobile drivers such as local road and traffic conditions. The voice interface could also be used with a minimum of distraction to control various local functions in the vehicle such as radio channel scanning or cassette tape operation.

Natural Language Conversions with Computers

Conversational Hypertext Access Technology (CHAT) is a technology developed at the Communications Research Centre Behavioural Research Division which allows people to retrieve paragraphs of information about a specific topic by asking questions in English or French. Consecutive questions and answers resemble a conversation or dialogue with the computer.

When a person types a question in natural language, the software compares the question to a number of predesigned templates which are attached to the paragraphs of text. After the software determines which templates best match the question, it displays the text associated with those templates. Thus, expert knowledge about a topic can be represented by both the paragraphs of text and the design of the templates.

Research which presents this technology to users has demonstrated that a relatively small number of templates can provide access to the answers to most people's questions. As well, it has shown that people are very tolerant of occasional replies that do not directly answer their questions. Ongoing research is directed toward automating input of the information and testing CHAT with larger databases and a variety of topics.

For further information please contact:

Dorothy Phillips, Ph.D.
Behavioural Research
Communications Research Centre
3701 Carling Avenue
Ottawa, Ontario
K2H 8S2

Department of National Defence

The Canadian Department of National Defence (DND) does not have a coordinated program of artificial intelligence research and development, nor does it have a coordinated program of implementing artificial intelligence technology into its operations. The Department treats AI as a technological area to be investigated as required to provide methods of solving particular classes of problems. As far as can be determined, the only fielded application in the Canadian Forces is a system to advise on maintenance procedures for a particular subsystem of the Automated Defence Digital Network. However, at the different Defence Research Establishments and the Operation Research and Analysis Establishment there are, and will continue to be, individual research efforts, primarily into the application of AI techniques, which may eventually lead to the introduction of such techniques into Canadian Forces operation.

The following is by no means exhaustive but gives a flavor of the types of AI applications that are being pursued or contemplated at the Defence Research Establishments. Some of the work is being done at the Establishments and some is being performed under contract.

Defence Research Establishment Atlantic

At the Defence Research Establishment Atlantic (DREA), Dartmouth, Nova Scotia, research and development has been ongoing for a number of years on the applicability of AI techniques to vessel detection and classification from passive sonar signals. The present work is investigating passive sonar feature analysis aids, automation of feature detection, and knowledge engineering for development of an expert system knowledge base and classification assistant. A concept system is being implemented and will provide a tool for refining and evaluating AI techniques.

Defence Research Establishment Valcartier

At the Defence Research Establishment Valcartier (DREV), near Quebec City, a project is underway to investigate how tools built using commercial expert system shells can assist a military analyst to perform intelligence functions that depend heavily on enemy order-of-battle and situation knowledge. These studies will eventually extend into related combat intelligence functions including templating and the prediction of probable enemy courses of action. A small investigation has been completed of manual and automated methods for acquiring knowledge from military experts. Another group is looking into the applicability of AI (expert systems, in particular) in the area of tactical information fusion, by developing an experimental prototype to study the feasibility of automating parts of the fusion process. A third group intends to apply expert systems to Threat Evaluation and Weapon Assignment (TEWA) problems in the maritime environment.

Defence Research Establishment Ottawa

The Defence Research Establishment Ottawa (DREO) is investigating the applicability and effectiveness of artificial intelligence techniques in a number of areas including satellite communications, communications network management, synthetic aperture radar, ship classification, single and multiple sensor tactical data fusion, situation analysis, weapons coordination, and pattern recognition. In satellite communications AI is being investigated for processing data related to stresses which may be encountered by such systems, and for deducing appropriate responses. In radar, research is underway to evaluate various algorithms to achieve automatic ship classification of high resolution radar images of ships. These algorithms, in general, encompass the fields of artificial intelligence, neural networks, and statistical analysis. With regards to artificial intelligence algorithms, the use of syntactic pattern recognition, for robust image feature recognition, and the use of the Dempster-Shafer technique, for the combining of independently made decisions, appear to provide viable solutions to the research problems.

Defence and Civil Institute of Environmental Medicine

The Defence and Civil Institute of Environmental Medicine (DCIEM), Downsview, Ontario, has interest in the use of AI techniques for robotics, computer vision, and machine learning, dynamic pattern recognition, human–computer interaction, and intelligent tutoring. In robotics the activities at DCIEM aim to advance the recognition of 3-D objects using stereoscopic vision, aiming toward semi-autonomous and autonomous operation of robots in unstructured environments. A computational vision system is under development for performing model-based object recognition (e.g., recognition of objects in aerial photographs of military targets). Brain modeling work on incentive learning is being done from a computational perspective which employs neural network representations and processes. Studies are conducted in the dynamics of tightly coupled genetically grown neural networks for the recognition of dynamic patterns such as speech. Interest is centered both on the network as a practical recognizer and on its role as a model for dynamic theories of human cognition. In the area of human–computer interaction, a framework called Layered Protocols and Explicit Models has been developed, based on both nonlinear dynamics, coding theory, and AI modeling. Each communicating partner models the other, including the other's models of the task, world, dialogue, and partner. The framework is being used in the development of quasi-intelligent interfaces. In intelligent tutoring, there are two projects related to AI. One involves the use of Kelly's personal construct theory to develop a model of learning that can be used in an intelligent tutoring system for training soldiers to recognize visual targets. The other involves using an expert system shell to develop an Intelligent Instructor's Aide to compare a student's performance on a gunnery simulator with PIECES 2, a model of the expected behavior.

Defence Research Establishment Suffield

The Defence Research Establishment Suffield (DRES), Ralston, Alberta, has no current project which uses explicit AI techniques. However, efforts on stand-off minefield detection are expected to involve expert systems in the near future and neural networks are being explored for unmanned airborne vehicle autopilots. Work is underway to develop a teleoperated robotic test bed for use in a number of roles such as ordnance detection, runway damage assessment, and reconnaissance, with plans to investigate increased vehicle autonomy in the future.

Defence Research Establishment Pacific

The Defence Research Establishment Pacific (DREP), Esquimalt, British Columbia, is interested in the application of artificial intelligence techniques to image processing and signal detection problems. AI is also being considered for applications such as aids in the analysis of sonar records for mine countermeasures.

Operational Research and Analysis Establishment

The Operational Research and Analysis Establishment (ORAE), Ottawa, Ontario, is investigating the use of artificial intelligence techniques in a wide variety of areas as budget planning and scheduling, logistics model, emergency response, search and rescue, battle management, and radar control and track initiation. ORAE is also interested in the field of neural networks for classification and optimization problems. In budget planning and scheduling, a prototype expert system was developed that makes some of the adjustments to the current army program of capital equipment acquisitioned taking into account changing priorities, budget restraint, and many other hard and soft objectives. This is part of a larger intelligent database decision support modeling system. A study investigating the use of an expert assistant to advise on which logical model to use in a given situation is underway. A prototype expert system for advising emergency response managers on how to deal with chemical spills has been developed. In search and rescue an expert system aimed at incorporating the knowledge and decisions made be a search master is under development.

In addition, the use of neural networks for classifying sonar signals and for choosing a firing doctrine have been studied. Simulated annealing is being considered as a methodology for finding optimum flight routes for several aircraft in order to overfly a multitude of sites.

As AI techniques become more widespread and mature, it is to be expected that more AI technology will find its way into defence research applications and eventually into operational situations. The Establishments will continue to monitor the technology, apply it where appropriate, and, if needed, initiate research to address specific deficiencies of defence concern.

For further information, please contact:

V.K. Taylor
Directorate Research & Development
Communications and Space
National Defence
National Defence Headquarters
Ottawa, Ontario
K1A 0K2

Energy, Mines and Resources Canada

Energy, Mines and Resources Canada (EMR) conducts research and development in artificial intelligence in several sectors of the department. In this report, we briefly describe the activities related to artificial intelligence at the Canada Centre for Remote Sensing (CCRS), the Canada Centre for Mapping (CCM), the Canada Centre for Mineral and Energy Technology (CANMET), and the Geological Survey of Canada (GSC).

Canada Centre for Remote Sensing (CCRS)

AI Activities at CCRS

Remote sensing is a technology which faces increasingly complex data fusion problems which require specialized expertise and experience to resolve. Artificial intelligence is an essential technology to capture the expertise as it evolves and to automate the procedures required for solving these problems and extracting meaningful information. Simple experts can be used to drive existing algorithms and interface the results into higher level "smarter" experts. The higher level experts may, in turn, be used to derive an analysis plan based on the user's goals. For example, a user who wishes to use remote sensing data for inventory map update will require knowledge from many disciplines in order to produce a meaningful and accurate final product. In the field of remote sensing, there are many spaceborne and airborne platforms available with a variety of sensors. The incorporation of these necessary components of knowledge from several disciplines in a modular fashion can be accommodated by a system of distributed advising expert systems, each having a specific scope of expertise.

The Canada Centre for Remote Sensing has conducted research for more than 8 years in artificial intelligence and for more than 18 years in computer vision as applied to the extraction of information from remotely sensed data. In 1983 a Prolog-based shell (Remote Sensing Shell, RESHELL) was created which was specifically designed for the interpretation and analysis of remotely sensed data. Figure 2 shows the elements of RESHELL. With RESHELL, expert systems can be created to execute processes in other languages on distributed computers, yet present to the user a seamless integration of diverse functions. This capability led to the creation of the first prototype for a hierarchical expert system, the Analyst Advisor, which can help nonexperts to correctly analyze imagery from the LANDSAT Thematic Mapper. The research and development was executed by a team of government, university, and industrial scientists. CCRS also develops single expert systems for specific tasks, such as cloud detection and snow mapping.

Computing Environment

The CCRS presently develops expert systems on MicroVAXes in Prolog. RESHELL is now being ported to Unix. These expert systems control processes and devices on other larger computers to achieve the multisource data integration. The image analysis processes involve more than 1.5 million lines of Fortran code. Devices include traditional computer peripherals, X-Window terminals, image displays, a Sun 490 compute-server, geographic information systems, and erasable optical disks. Macintosh computers are also used at CCRS as AI platforms.

RESHELL

RESHELL (Figure 2) uses meta rules and object rules. Meta rules provide the decision criteria for groups of rules (phases) in the expert system, while object rules are used to prove or deduce object values using either forward chaining or backward chaining, and have an associated uncertainty. Meta rules are grouped in phases, in

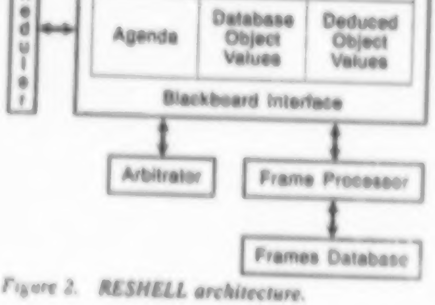

Figure 2. RESHELL architecture.

which all the rules in a certain phase (the current phase) are tested. The actions of a meta rule may set the next phase. This reduces the search space while testing the conditions of the rules and allows for a logical order of operation of the expert systems. Object rules provide the data as criteria for meta level decisions; thus they are invoked by either conditions or actions from the meta rules. The consultation between experts is performed at the meta level, which determines the hierarchy of the system as a whole. Each expert has its own blackboard in which it may paste objects or data for use by other rules at either the meta or object level. Data passed between experts can also be through a frames database, which stores objects as frames.

A specialized facility has been developed as a part of the shell for automated program execution. While the initial intention of the facility was for the operation of image analysis programs, there are few restrictions on the type of program or procedure which may be controlled. The automated program execution facility is entitled the Interactive LDIAS Task Interface (ILTI). The ILTI is implemented as an expert in RESHELL with specialized functionality. Each LDIAS program which requires execution by the expert system must have an ILTI with knowledge about its operation.

Given the complexities of combined systems for image analysis and geographic information, we believe that window-based expert systems can simplify the user interface and reduce the educational level required to effectively operate these systems. The expert shell (RESHELL) at CCRS has a DECwindows interface implemented at the level of the shell (in Prolog). Specific functions are performed automatically at the Prolog level, such as window sizing, widget selection, error messages, and string-to-Prolog translation. A comprehensive rule and object editor has been implemented as part of this DECwindows interface.

The CCRS is currently implementing this user interface in its expert systems built using RESHELL, and conforming the entire user interface to MOTIF standards. Included with this development is adapting the user interface to the expertise level of the user, cut–copy–paste operations, dynamically building multipurpose prompt windows within a parallel-processing, nondeterministic prompting environment, window logging, and a design for the easy addition of future developments in the expert systems.

The MOTIF standard user interface will facilitate porting the interface along with the shell to other operating systems (VMS, ULTRIX, POSIX, Unix, and MacOS). Ultimately, we envision a multivendor system incorporating large and microcomputers which will be distributed over a large geographic area. Industry is involved in this development and we anticipate RESHELL being made available in the first quarter of 1991.

Analyst Advisor

Our first hierarchical expert system created was called the Analyst Advisor. This hierarchical knowledge-based expert system guided a user through a session on the LANDSAT Digital Image Analysis System (LDIAS) by parsing and operating programs which require distributed knowledge and high expertise. The integration process is complex and is highly data dependent. Data accuracy and processing recommendations were reported at each step of the process. The user made decisions based on this information relative to previously set goals for classification accuracy and map update requirements. The knowledge for the processing sequences and rule base for the decision making were coded in Prolog. The high-level knowledge or "meta rules" control initiated and directed Fortran tasks to perform the lower level image processing. The Fortran tasks ran on a number of VAX nodes in a local area cluster environment.

System of Hierarchical Experts for Resource Inventories (SHERI)

The CCRS is developing a second hierarchical expert system, called System of Hierarchical Experts for Resource Inventories (SHERI). The SHERI Project was approved in October 1989 through a competitive process within the Government of Canada administered by Industry, Science and Technology Canada (ISTC). The SHERI Project has as its goal the creation of an intelligent system for updating the forest inventory of British Columbia (more than 6000 1:20 000 maps or GIS files) using Thematic Mapper (or SPOT) imagery. SHERI will contain more than 150 expert systems. The project will be completed by March 31, 1993. Two federal departments, Energy, Mines and Resources Canada and ISTC, and the British Columbia Ministry of Forests are supporting the project. The SHERI Project is also supporting the commercialization of our expert systems shell, RESHELL.

As the second largest country in the world, Canada has a strong domestic need for indigenous technologies which support geographic information systems and resource management systems. Because the information in resource inventorying systems becomes obsolete quickly, it is essential that remote sensing data be used to update these systems frequently. Past methods of information gathering for forest inventorying have involved extensive use of aerial photography. Because of the costs of photo acquisition, the lack of digital products, and the variety amongst photo interpretations, it has not been possible to update the forest inventories on an annual basis. It has been demonstrated that remote sensing is an effective and economic tool for such updates. To use the remote sensing tool requires staff with rare expertise in digital analysis of remotely sensed data and the integration of the interpretations with the geographic information systems. The complexities of the analysis methods and information systems can be handled by knowledge-based methods. What is needed are expert systems for computer vision of remotely sensed images and the integration of such vision expert systems with existing resource inventorying systems.

Within the SHERI project, CCRS conducts research and development on a hierarchical expert system for integrating remote sensing data from multiple platforms and sensors with a variety of geographic information systems in order to inventory forests and advise on forest management strategies over time. This project will benefit the management and exploitation of one of our most important renewable resource, forests. The outputs of the research will include the implementation of a forest resource inventorying system using expert systems, the commercialization of a shell with unique capabilities, publications, and training. This forest resource inventorying system will increase the usage of remote sensing data, reduce the requirements for staff with rare skills, and ensure consistent accuracies for the provincial resource database. The system will also serve as a useful tool to train existing operators in the analysis of remotely sensed imagery in order to recognize forest objects of interest, such as harvested areas, areas of poor growth or poor regeneration, and areas not conforming to approved timber cutting profiles. Current capabilities within the system are data preparation, geometric correction, terrain correction, image classification, segmentation, cloud detection, raster-to-vector and vector-to-raster conversion, and GIS map overlay with remote sensing data, filtering, and vectorization of the final products. Many advanced artificial intelligence techniques are being used in SHER!, including automated program execution, uncertainty propagation, and machine learning.

Photo-interpretation Keys Expert System (PIKES)

For mapping, creating resource inventories, and environmental monitoring, one makes use of imagery acquired from aircraft and space platforms. The process of a human interpreting this imagery is known as photo interpretation. There is a shortage of photo interpreters. Photo interpreters presently use visual keys from a set of small visual examples of various features that can be identified on aerial photos to help them identify and label features. However, these keys are limited in scope, frequently of poor design, and are not updated often enough to take advantage of new sensor technology. Less-experienced photo interpreters often show a variability in their interpretations indicative of the improper development of photo guides or "keys." A large potential for improvement in photo interpretation exists if the interpreter can use remotely sensed digital images and be guided in the creation of keys with greater scope and robustness.

This project will develop an expert system to train photo interpreters with the help of on-line visual keys. It has also been recently approved for funding by the AI R&D Fund of Industry, Science and Technology Canada.

The use of an expert system to build image interpretation keys and train photo interpreters with these keys offers unique capabilities not available from other computational tools: expert knowledge can be easily updated without reprogramming; the training can be, to a large extent, self-paced and requires a minimum of supervision; a higher level of accuracy and consistency in individual interpretation will be achieved by the common training base. This expert system will capture the knowledge of photo interpreters for decisions made in both key selection and in the interpretation process. Operation of the expert system will invoke decisions for photo interpretation in a similar pattern to the human experts; from initial separation of land-use classes to species and age determination of vegetative and forest cover. During the key selection process, image portions can be saved and archived along with a database index, which will be activated during future photo-interpretation procedures when the particular class type is suggested to the user. The keys will be used during photo interpretation of digital satellite or aerial images. The impacts of the project will be increased used of digital remote sensing imagery, less costly training of photo interpreters, the availability of distributed training, the improvement in accuracies of photo interpretation (with impacts on mapping and monitoring), and the export sales of low-cost systems to countries without expertise in photo interpretation.

Environmental Monitoring System (EMSys)

The evolution of the science of remote sensing has led to significant advances in the instruments of the science and in the cognizance of its potential use for large-scale and long-term monitoring of the Earth. Several international interagency studies are currently being initiated for long-term monitoring using advanced remote sensing tools specifically designed for global monitoring, including spaceborne imaging spectrometry and radar sensors. The Canada Centre for Remote Sensing is participating in several studies of global monitoring.

While global monitoring will continue for many years, the science required for the monitoring is developing now, and the development of an intelligent tool for incorporating the science from many domains is essential. The problem is akin to creating a national geographic database through the use of intelligent agents to enable data sharing and processor sharing for analysis and decision making. The use of expert systems for incorporating and analyzing data in the context of the monitoring studies will enable monitoring problems of continental scale to be addressed and the intelligence in the tools will ensure their adaptability.

The CCRS has proposed conducting research and development on a hierarchical expert system (EMSys) for monitoring and understanding changes in the boreal forest biome and arctic land areas. To increase our understanding of the interactions between boreal forests and atmosphere and the responses of the boreal forest to climatic warming, NASA and CCRS will be conducting a 5-year experiment, the Boreal Forest Experiment, on two sites in boreal regions within Canada, which are among the most sensitive areas to global warming in the world. In addition, CCRS and the British Columbia Ministry of Forests are presently conducting joint research in mountainous areas sensitive to environmental changes. The collected data will need to be integrated with existing Canadian geographic information, digital terrain models, forest growth models, climatic model outputs, atmospheric correction models, and environmental monitoring models. These experiments serve as important sources of knowledge acquisition for the expert systems to be built in EMSys. Experts in EMR, the British Columbia Government (Forests, Crown Lands, and Environment), Environment Canada, and Forestry Canada will also be used to provide knowledge to EMSys. The outputs of the research would include the development of the first prototype system for environmental monitoring over continental scales using AI technology, the testing of the system with a complex data set involving new remote sensing measurement techniques such as imaging spectrometers and polarization radars, the implementation of a system which can integrate geographic information from federal, provincial, and foreign governments, methodology to evaluate the information value of multiple data sources, the ability to use distributed, parallel expert systems, commercialization of the prototype with unique capabilities, publications, and training.

The objectives of large-scale monitoring are well defined; however, the expertise required is widely distributed and the analysis functions involved are highly complex and, in some cases, still under development. The monitoring activities underway in this area are currently distributed among highly specialized, independent groups. This project intends to integrate the experience and knowledge of these experts through the integration of this expertise in an artificial intelligence, decision making, and prediction system.

System of Experts for Intelligent Data Management (SEIDAM)

SEIDAM is a project proposal to conduct research on a system of expert systems for intelligent data management. The CCRS has much expertise in developing systems for integrating geographic information with space and aircraft remote sensing data and in managing large archives of remotely sensed data. SEIDAM will be composed of expert systems grouped in three levels. At the lowest level, the expert systems will manage and integrate data from diverse sources, taking account of symbolic representation differences and varying accuracies. Existing software can be controlled by these expert systems, without rewriting existing software into an AI language. At the second level, SEIDAM will take the interpreted data (symbolic and numerical) and combine these with existing data models. At the top level, SEIDAM will respond to user goals for predictive outcomes given existing data. The SEIDAM Project will address the research areas of expert systems, data management, storage and retrieval, and user access and interfaces.

Involvement of the Private Sector and Universities

Following CCRS tradition, all its AI projects are done in close collaboration with industry and research groups in different universities. For example, the SHERI project has contracts with Intera-Kenting Technologies, Price Waterhouse, the University of Ottawa, Ferihill Technologies Ltd. of Victoria, and the University of Toronto.

Contact Person

For further information on the CCRS AI activities, please contact:

Dr. David G. Goodenough
Canada Centre for Remote Sensing
Energy, Mines and Resources Canada
1547 Merivale Road
Ottawa, Ontario
K1A 0Y7
Telephone: (613) 952-2760

Canada Centre for Mapping (CCM)

Rule-based Generalization for Context Transformations

Background

In thematic mapping, scale reduction is an important part of map production. The production process usually includes the transfer of detailed information from larger-scale maps to maps of a smaller scale. Automated generalization, like manual generalization, revolves around the following activities: simplification, displacement, exaggeration, combination, selection, classification, and symbolization. Although these activities are applied to both the graphic elements and thematic content of the maps, their involvement is not equally distributed. For example, simplification, displacement, and exaggeration are tools directed towards the manipulation of the graphic presentation of spatial objects, and their intervention is strongly, but not exclusively, motivated by legibility requirements. On the other hand, selection and classification affect primarily the thematic content of maps, and their involvement is prompted by the type of information to be displayed. Hence, the former activities emphasize form, whereas the latter emphasize content. However, those classes of activities are strongly connected, since selection and classification lead necessarily towards changes in graphical presentation and the needs for graphic transcription may, in turn, promote further selection and classification.

Previous efforts in automated generalization concentrated mainly in the graphical manipulation of object representation. Those manipulations were prompted by the need to maintain map legibility as the process of scale reduction. Corresponding criteria to judge the performance of cartographic generalization methods were exclusively geometrical.

Research Objectives

The objectives in developing the rule base are to develop a more comprehensive approach to generalization in the digital environment and to address the contextual relationships among objects. The rules provide guidelines of when to generalize, what to generalize, and why. The automated process of contextual generalization involves the application of "necessity factors" in conjunction with a GIS database for thematic mapping. The process includes a set of rules that supports information requirements of base maps for presentation at different scales and maintains the relationships of base map elements to different thematic subjects. The rule base will include a subset of rules based on attribute information of objects that can be used as criteria for generalization. A subset of rules will be included to ensure that the integrity of spatial relationships is preserved through scale change.

System Environment

The basic computing environment consists of: Sun Workstation, Unix, and the Electronic Atlas.

The geographic information database will be created using ARC/INFO Oracle. The rule base is being tested using either Brainmaker or Smalltalk, with results implemented in an Electronic Atlas mapping package within the Canada Centre for Mapping.

Contact Person

For further information, please contact:

D.E. Richardson
Technology Section
Topographical Mapping Division
Canada Centre for Mapping
Energy, Mines and Resources Canada
615 Booth Street
Ottawa, Ontario
K1A 0E9
Telephone: (613) 992-4683

Statistical Mapping Advisor

Background

Statistical mapping is drawing more and more attention from cartographers for thematic mapping due to its relatively high mathematical accuracy, and due to its provision of data for synthetic mapping. The most commonly used synthetic mapping technique is classification, which, for example, classifies a population into several groups so that the spatially distributed phenomena can be represented on maps in class interval series. In addition to this, however, there are some more sophisticated statistical analysis techniques, such as clustering and discriminate analysis, that are useful for thematic mapping, but have only infrequently been used. With the development of GIS software, however, multivariable synthetic mapping became feasible and necessary. Accordingly, applications for sophisticated statistical analysis techniques are presently being explored at the Canada Centre for Mapping.

A prototype expert system, Statistical Mapping Advisor (SMA), is being developed to act as an expert in statistical mapping and to provide users with advice while using statistical analysis mapping techniques. To use SMA no prior knowledge of statistical analysis is required. By consulting the system, users can expect advice on what the related statistical analysis methods mean, what types of maps and data they can be used for, what methods should be adopted for specific applications, how these methods should be used, and how they can be visually displayed.

System Environment

The basic computing environment is Sun Workstation, Unix, and MS-DOS.

The Statistical Mapping Advisor (SMA) is designed for applications for the National Atlas Information Service and is being implemented on a Sun workstation. The statistical analysis component is being tested using STATgraphics which are compatible with the Sun workstation and supported by MS-DOS.

Contact Person

For further information, please contact:

Dr. E.M. Siekierska
Engineering and Systems Development
Geographical Services Division
Canada Centre for Mapping
Energy, Mines and Resources Canada
615 Booth Street
Ottawa, Ontario
K1A 0E9
Telephone: (613) 992-4470

Investigation of Automated Name Placement for Topographic Map Production

Background

Automated name placement in topographic mapping has remained a difficult problem due to the inherent aesthetic component of placing names. A basic set of rules for name placement exists (e.g., names should not overlap each other, should be placed horizontally if possible, and should not overlap or conceal other map features), but the final placement is also guided by what looks good from a cartographic standpoint. Several solutions have been developed for automatically placing names. All of them solve some part of the automated name placement problem (normally broken into the subproblems of point, linear, and area feature name placement), but none have gone beyond the research prototype stage. Of particular importance is (i) the failure of most existing software to consider the structure and completeness of existing map and name databases, (ii) the absence of information from experienced cartographers to establish a more complete rule set for name placement, and (iii) the lack of consideration of implementing a computer-assisted approach when the fully automated approach fails.

Research Objectives

This research has the following objectives:

- Review the literature on automated name placement and compare these automated results with manually placed names on Canadian topographic maps.
- Develop a rule base for name placement on Canadian topographic maps based on interviews with experienced cartographers.
- Develop algorithms for automated placement of names consistent with the data structure of the National Geographic Names Data Base (NGNDB) and the National Topographic Data Base (NTDB).
- Implement these algorithms and test them using sample digital data from the NTDB, along with names in the NGNDB for the same area.
- Outline the steps required to implement automated or computer-assisted name placement for topographic map production at the Canada Centre for Mapping.

Contact Persons

For further information, please contact:

Mr. D. Brassard (Research Scientist (EMR Research Agreements Program))
Technology Section
Topographical Mapping Division
Canada Centre for Mapping
Energy, Mines and Resources Canada
615 Booth Street
Ottawa, Ontario
K1A 0E9

Dr. Bradford Nickerson
Assistant Professor
School of Computer Science
University of New Brunswick
P.O. Box 4400
Fredericton, New Brunswick
E3B 5W9

Canada Centre for Mineral and Energy Technology (CANMET)

As a designated federal government technology center, the Canada Centre for Mineral and Energy Technology (CANMET) conducts research and develops technology for the Canadian minerals, metals, and energy industries. CANMET's objectives include: to help these industries become more competitive; to improve health, safety and the protection of the environment within the industry; to diversify and improve the efficient use of mineral and energy resources; and to support government policy initiatives.

CANMET's mission is to assist the mineral, metal, and energy industries and as such it does not do research in artificial intelligence as a science; rather, CANMET employs artificial intelligence in its work where feasible.

The minerals industry has traditionally improved its competitive position by reducing staff. Most companies are now considering other cost-saving measures such as automation and process control. CANMET's clients are increasingly making use of sensors, meters, gauges, and computers to monitor, record, and instantly analyze process conditions. Expert systems will be key to automating the operation of mineral processing plants. CANMET has expertise in using expert systems to design and operate ore processing plants, and will be applying this expertise in energy processing.

CANMET's efforts in artificial intelligence have been directed mainly towards developing the knowledge base, mathematical models, and computer simulation techniques necessary to form the basis of an expert system. CANMET has participated in the development of several prototype expert systems and is currently working towards the development of a number of others which will be described below. The systems currently being developed will be customized to existing research facilities, and CANMET will work with client companies to modify the systems for specific commercial applications.

Mineral Science Laboratories (MSL)

Introduction

CANMET has completed three projects on expert system prototyping in mineral processing between 1987 and 1989. Current conclusions on the pros and cons of expert systems as perceived through these projects are excerpted from the paper "Latest Developments in Computer Aided Mineral Processing Techniques at CANMET" presented at the APCOMIN' 90, Application of Computers in the Mining Industry, Santiago, Chile, May 14–17, 1990, CANMET division report MSL 90-41 (OP&J).

Expert systems have been available to applied R&D users since the mid-80's. They were often promoted as a revolutionary technology. After five years and many prototypes, it is fair to say that the expert systems shining star has somehow lost its initial brightness. At least four factors have contributed to the apparent fading of the technology:

1. the quick absorption of new technologies into the main-stream by the scientific and engineering communities may be, in part, responsible for the decreased visibility of expert systems technology by making it a textbook item
2. the quick pace of technological innovation in computing contributes to the rapid obsolescence of yesterday's novelties (e.g., LISP machines)
3. the overexpectations laid on expert systems in parallel with an underestimation of some difficult steps in their development and integration into existing software, especially on-line, may have discouraged some early converts to apply expert systems in their operations
4. the overall abuse of the terminology has diluted the original concept and contributed to some disillusionments and a drop in popularity of the name "expert system"

Figure 3 shows the various roles an expert system may play in a mineral processing environment: diagnosis (trouble shooting), prognosis (operator adviser) and on-line contribution to process control. Over the last few years, CANMET has experimented with expert systems prototypes in all three applications as briefly summarized below (Laguitton, 1989b).

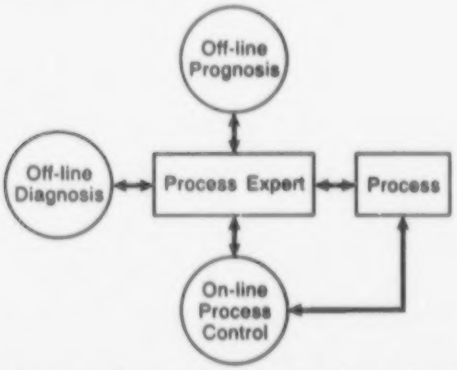

Figure 3. Process expert system on-line and off-line configurations.

Off-line Expert Systems

The first off-line system sponsored by CANMET was a diagnosis system. The clinker grinding expert system developed with Canada Cement Lafarge (Flament, 1987; Flament, 1988) simulated the reasoning of a specialist whose chief responsibility was to resolve quality and production level failures in clinker grinding circuits. It operated in the KEE/Explorer environment. A second expert system prototype was developed during the same period by a team from Lavalin and Canadian Electrolytic Zinc (CEZ), for operator assistance in the control of a zinc leaching and refining plant (Laguitton, 1989c). Process prognosis was therefore the objective of this system. The use of conventional hardware was the chief difference from the clinker application described above. The integration of the system within a conventional computer architecture was of prime concern from the onset and the TI Personal Consultant shell was selected with a COMPAQ-386 which provided adequate speed for the inference with a 250-rule knowledge base.

On-line Expert Systems

On-line expert systems reported so far in mineral processing are preciously few. Such an application of expert systems in automatic process control was conducted jointly by CANMET and Kidd Creek Mines (Burrows, 1989). It consists of a data validation and substitution expert system integrated at the interface of an X-ray analyzer and an existing automatic process control and plant information computer. Failures or pre-failures in the automatic sampling and analysis system are identified by consultation of a knowledge base in the Nexpert RUNTIME shell supplied by Neuron data. Assay substitutions are performed by various algorithms, using reliable historical process data and numerical methods. The validated process data are then cleared for the automatic process control system and/or the plant information displays. Actual on-line operation testing of the system installed on a MicroVAX computer was conducted, and the "substituted" values were displayed on the operator's stations for several months. However, it was soon discovered that a significant number of the "suspect" assays were actually due to process fluctuations rather than sampling and/or analytical failures. Work is under progress at CANMET to take this into consideration to allow process fluctuations to be handled in dynamic simulation algorithms and improve the substitution technique for "truly" invalid assays.

Conclusions

The following are the main conclusions drawn from three years of expert systems prototyping at CANMET:

1. Expert systems are often needed in the context of existing computer hardware and software; integration is therefore of the outmost importance and must be considered closely during the planning phase.
2. Expert systems were initially promoted as having been designed to handle "heuristic" data, often with symbolic (e.g., words) rather than numerical variables. However, integration with other software often requires conversion from symbolic to numeric, calling for extra expertise which may not be available and, to a certain extent, destroying the original concept of "symbolic" knowledge. Furthermore, few expert systems shells include the capability of calculating certainty factors during the rule inference. The old adage has been revisited in the context of expert systems: "a formula is worth a thousand rules" (author unknown).
3. The development of expert systems has some definitive advantages in forcing the developers to conduct a thorough inventory of their knowledge and to establish priorities among the many problems potentially addressable by expert systems.
4. Expert systems, as introduced in the 80's, seem to have been only the first generation of the software engineering tools of the 90's. Virtually any software application does and will offer characteristics of both numerical algorithms and symbolic rule-based inference and it will be increasingly difficult to precisely use the name "expert system" in its original meaning to describe them.
5. A great deal of progress is still needed in developing and teaching efficient knowledge engineering techniques. This observation may be another way to describe the rare blend of rational thinking and interpersonal communication that should be produced by a quality scientific education.

References

Burrows, M.J., Carriere, J.D., Leung, J., Laguitton, D. (1989). An Expert System Designed to Improve Reliability of Assay Information in a Flotation Plant, 21st Canadian Mineral Processors Conference, Ottawa, Ontario, Canada, Jan. 1989.

Flament, F., Vanderstichelen, J., Laguitton, D., Hodouin, D. (1987). Prototyping of an Expert System for Trouble Shooting of Clinker Grinding Mills, presented at the Seminar on Artificial Intelligence in Minerals and Materials Technology, The University of Alabama, Tuscaloosa, Oct. 1987.

Flament, F., Vanderstichelen, J., Laguitton, D., Hodouin, D. (1988). Tool Assessment of Expert System Prototyping: the KEE Expert System Shell, *Proceedings of the First Canadian Conference on Computer Applications in the Mineral Industry*, Laval University, Quebec, Canada, March, 1988.

Laguitton, D., Leung. J. (1989b). Advances in Expert System Applications in Mineral Processing, 28th Annual CIM Conference of Metallurgist, Halifax, Nova Scotia, Canada, Aug. 20-24, 1989.

Laguitton, D., Bilodeau, M., Ghibu, C., Boulanger, J.P., Farner, I., Dupuis, D. (1989c). Expert System Adviser for Operators of a Zinc Leaching Plant, *Proceedings of the 6th Symposium on Automation in Mining Mineral and Metal Processing*, IFAC pub. Pergamon, Oxford, Buenos Aires, September, 1989.

Contact Persons

For further information, please contact:

Dr. Daniel Laguitton / Jean Leung
CANMET
Energy, Mines and Resources Canada
552 Booth Street
Ottawa, Ontario
K1A 0G1
Telephone: (613) 996-7953/992-633
Fax: (613) 954-1563

Mining Research Laboratories (MRL)

The Mining Research Laboratories (MRL) are working towards the development of systems to be used by engineers in the design of ground support systems and the assessment of ground control. The knowledge-based system will combine various mine planning modules and mathematical models for rock stress analysis into an integrated package for complete analysis of mine design and safety.

Metals Technology Laboratories (MTL)

The Metals Technology Laboratories (MTL) are working towards expert systems in three technology areas:

- A system is being developed for the prediction of residual fatigue life for welded structures. This system will provide the methodology for predicting the increased chance of failure or the extrapolated residual life of a structure which has sustained fatigue cracks.

- The Foundry Technology expert system will guide a casting designer/purchaser to the best choice of alloy and manufacturing process for a specified application. MTL has already developed the sensor array and expert system to control costs of electric resistance furnace operation. Other systems under development include mould design and preparation linked to NC machining methods; design of gating, running and risering of the mould cavity and pattern production; foundry sand quality control; and an expert system for safe sand recycling or disposal.

- The third area MTL is working in deals with the control of the Hot Strip Rolling of Metals. The initial system will capture data from rolling research and then be capable of predicting a detailed rolling schedule which when followed would produce a rolled alloy with the desired properties. Ultimately the system may be expanded to control commercial rolling mills.

Coal Research Laboratories (CRL)

The Coal Research Laboratories (CRL) are developing a prototype expert system for process control of the coal fines recovery section of a coal preparation plant. The system will be initially developed to control the pilot scale plant at CRL and will be expanded to commercial scale after it has been validated.

Energy Research Laboratories

The Energy Research Laboratories are developing a program which will correlate the behavior of a coal in a number of coal gasification processes to the physical and chemical properties of the coal. The aim of the program is to create an expert system for assessing the suitability of any given coal in a gasification process based on the physical and chemical properties alone, without having to run a gasification test.

Efficiency and Alternative Energy Technology Branch

The Efficiency and Alternative Energy Technology Branch is active in developing three expert systems:

- The Diagnostics and Training for Building Energy Management Systems expert system will help improve the utilization of building heating, ventilation and air conditioning systems by providing techniques for training building operators and also diagnostic tools to assist in identifying problems.

- The Large Building Energy Efficient Design expert system is intended to allow the design professional to understand and evaluate the range of options available to make informed energy related design decisions.

- The Enhanced Energy Simulation for Low Rise Buildings expert system will provide rules, procedures, and guidance on key issues or concerns that must be addressed at the design stage to avoid post-construction problems. These problems relate to the energy performance and the indoor air quality and safety of buildings. The system will be based on the HOT2000 Energy Analysis Program which was developed by EMR.

Geological Survey of Canada

Introduction

At the present time activities at the Geological Survey of Canada (GSC) associated with artificial intelligence (AI) are being driven from the bottom up by a few geoscientists who believe that AI techniques have a role to play in meeting the objectives of their particular projects. Currently there are no top-down projects where objectives and resources have been specifically defined and committed to AI, by Sector Management. Therefore, annual expenditures on AI in the Geological Survey are of the order of 10K$ (excluding salaries), and involve less than 10 project officers.

Summary of Activities

Geological Survey interest in the various fields of AI is largely confined to production rule expert systems (ES). One project sees an application of neural networks (NN), which will be investigated by years' end. In the following paragraphs activity in the various Divisions of the Sector are briefly described. This information is also summarized in Table 5.

The most active group in attempting to apply AI techniques is in the Mineral Resources Division, Ottawa, where four projects have AI components. Three of these are closely related to mating expert system methodologies to Geographic Information Systems (GIS) to undertake mineral resource assessment tasks. The work is microcomputer based and relies upon the "unique conditions" tables from within the SPANS (TYDAC

Table 5. AI projects at the Geological Survey.

Division	Contact Person	Project
Atlantic Geoscience Center	B. Loncarevic	A hypertext-based technical manual has been prepared with VP-Expert 2.0 to provide advice on the use of shipboard oceanographic equipment.
Cordilleran Division	G. Woodsworth	Working as a domain expert with David Lefebure (B.C. Geological Survey Division) and Alan Campbell (Consultant, BC-GSD funded) on an ES to assist in mineral resource assessment in the North Coast Range, B.C.
Geophysics Division	R. Coles	Investigation of ES methodology to capture the knowledge and procedures concerning geomagnetic storm forecasting in order to improve the quality and timeliness of such forecasts.
Lithosphere and Canadian Shield Division	M. Schau	Several experiments have been made using free-ware from the AI community into the classification of metamorphic rocks and conflict resolution between competing classifications..
Mineral Resources Division	F.P. Agterberg	Working with G.F. Bonham-Carter on the application of the weights-of-evidence methodology to mineral resource appraisals via a GIS. Additionally, Convenor of the "Artificial Intelligence Modelling for Mineral Resource Evaluation" session of the IAGOD Symposium, Ottawa, August 1990.
Mineral Resources Division	G. Bonham-Carter	In addition to the work described above, machine-learning tools, e.g., KnowledgeSeeker, are being used to obtain rules for recognizing areas of high mineral potential directly from a GIS database. In addition, investigations are being made on forward chaining through rules concerning mineral deposit classification and localization, found in such ES as Prospector, directly on a GIS database in order to assess mineral potential.
Mineral Resources Division	C.F. Chung	Dr. Chung is investigating the role of knowledge representation in rule formulation and the impact of the choice of probability or certainty formulation, e.g., Bayesian vs. Dempster-Shafer, on inference outcome in connection with mineral resource assessment tasks using GIS databases. This work is to be presented at the "International Symposium in Mineral Exploration: The Use of Artificial Intelligence,," Tsukuba, Japan, October 1990.
Mineral Resources Division	R.G. Garrett	Developing an ES, ®GXpert:Au, using KnowledgePro with a retired member of staff, Dr. R.W. Boyle, who is a world expert on the geochemical exploration for gold deposits. This ES will be distributed as a publication of the GSC. Also investigating the application of Neural Networks to the classification of geochemical data in order to recognize areas of increased mineral potential. Additionally, Convenor of the "Intelligent Information Management: Knowledge-Based Systyems" session of the GEOINFO IV symposium, Ottawa, June 1990.

Technologies, Inc.) GIS for "data" input to be viewed against the "knowledge base." Currently the "knowledge base" is derived internally from the distribution of known mineral occurrences in the study areas. A variety of ad hoc and structured procedures for undertaking this task are being employed. In this context, the Ottawa-developed KnowledgeSeeker machine-learning software is being used. Investigations have recently started into using forward chaining models in conjunction with externally established production rules concerning the classification and localization of mineral deposits, such as those in the United States Geological Survey's Prospector ES.

The fourth project involves an ES, ©GXpert:Au, being developed as a geochemical survey planning advisor. This ES uses the KnowledgePro shell in conjunction with the over 40 years of experience, publications, and unpublished data holdings of a recently retired senior GSC scientist, Dr. R.W. Boyle. The objective is to provide users with recommended suites of geochemical pathfinder elements for gold deposits conditional on the deposit type or local geology, the sample media, and the surficial environmental conditions. The ES will be published as a GSC Paper in 1991 and will also test the receptivity of GSC's geochemical customers to this form of knowledge and technology transfer. Lastly, a Memorandum of Understanding (MOU) is being discussed between geochemists at the GSC and the British Geological Survey to jointly investigate the use of neural networks in classifying geochemical data in terms of its strength of indication of mineral potential.

Two other geoscientists in Ottawa, in the Geophysics and Lithosphere and Canadian Shield Divisions, have also been investigating the application of ES to the forecasting of geomagnetic storms and the classification of high-grade metamorphic rocks. The former application is more advanced. A Prolog prototype has been built and reviewed; the experiences gained from that exercise together with recent developments in forecasting philosophy will be incorporated into an updated ES. It is likely that a commercial shell will be employed in developing the ES. At the present time various shell options are being considered.

At the Atlantic Geoscience Center, Dartmouth, Nova Scotia, the hypertext functionality of VP-Expert 2.0 has been used to prepare a user manual for a variety of pieces of ship-borne oceanographic survey equipment. The application focuses more on the hypertext functionality than the knowledge base and associated branching decisions. However, the package has proven useful in the field in providing the appropriate information in a timely fashion to the degree of detail required.

Finally, one staff member in the Cordilleran Division, Vancouver, British Columbia, is providing domain specific geoscience knowledge to a British Columbia Geological Survey Branch AI project. This project is developing an ES to assist in mineral resource assessment in the North Coast Range mapping area. The AI component for this project is being provided by Alan Campbell of Campbell & Associates; Campbell was a member of the original Stanford Research Institute team that developed Prospector under contract to the U.S. Geological Survey.

Conclusions

Thus, it may be seen that at this time the Survey's involvement with AI is limited and is dominated by projects concerning mineral resource assessment. In the present human and financial resource environment it is unlikely that any major changes in level of involvement will occur in the near future. It is also unlikely that the Survey could afford any cost-shared projects with outside AI industry contractors, even 50% of the costs of such projects could cripple the remaining non-AI projects in a Division. However, limited development of the type described above will continue. Should the results of these projects be encouraging the case may be made for the commitment of new resources to AI-related projects. Additionally, Sector Management has recently shown an increased interest in AI, and if "corporate" AI applications are identified that could lead to major changes in the AI environment.

Prepared by D.G. Goodenough, J.-C. Deguise, R. Garrett, A. Lane, and D.E. Richardson.

Government Consulting Group

The Government Consulting Group (GCG) is the successor organization to the Bureau of Management Consulting (BMC). We are now part of a special operating agency, Consulting and Audit Canada. As previously, consulting services are offered to Government on a fully revenue-dependant basis.

The GCG continues the commitment and active interest that began in BMC in Strategic Technologies, including AI as a major component. For the last five years, we have successfully conducted a number of projects related to AI for our clients.

The Government Consulting Group supports its clients and provides opportunities for those in public and private sectors interested in AI to meet and exchange ideas. The GCG frequently acts through alliances among the private sector, universities, and government laboratories. By forming such alliances, it is able to provide the best resources wherever they are. Thus clients receive the best product possible and, at the same time, as much data and exposure as possible is fed back to the R&D community.

GCA Projects: AI and Related Disciplines

Automatic speech recognition (ASR)

- Phoneme recognition by neural net
- Automated system to evaluate ASR devices

Data fusion

- Neural network selection of firing doctrine for defence against missile attack

Sonar

- Signal recognition by neural network

War gaming and simulation

- Network of computers to game, simulate, and display naval warfare (platforms, strategy, tactics, map)

Expert systems

- Prototype design and specification for expert system to conduct toxicity assays — accept data, select, and perform life-table analyses

State-of-the-art surveys

- Canadian R&D of voice input technologies
- North American work in ICAI
- Canadian R&D (and use of) neural networks

Program development

- Coordination and support for inception of an AI program

Education and public awareness

- Introductory seminars, advanced tutorials, and publication in AI disciplines

For further information, please contact:

Zachary Jacobson
Government Consulting Group
365 Laurier Avenue West
Ottawa, Ontario
K1A 0S5
Telephone: (613) 995-9863
Fax: (613) 952-8295
net id: zack@ncs.dnd.ca

National Research Council

Introduction

At the National Research Council (NRC), the Division of Electrical Engineering Laboratory for Intelligent Systems (now the Knowledge Systems Laboratory, Institute for Information Technology) began its program in Applied Artificial Intelligence in early 1987. The activities of the program are directed towards the development of knowledge-based system tools and methodologies and their application in response to specific requirements. The general objective of the Laboratory is:

> To provide, through applied research and collaboration with Canadian industry and universities, intelligent systems technology which will make Canadian industry more competitive in the emerging fields of sensor-based robotics and knowledge-based systems.

Within this general objective, the Applied AI Program has the following specific objectives:

1. To establish a national laboratory capability in the field of applied artificial intelligence through the successful completion of projects which are technically challenging as well as having the potential for exploitation.
2. To establish a facility that will provide support for projects of NRC's Industrial Research Assistance Program.
3. To provide NRC with a facility for international collaboration in AI.
4. To collect data on AI technology and applications for easy dissemination to other groups within NRC and to Canadian industry.
5. To establish demonstration and training facilities for other groups within NRC and for Canadian industry.

Core Support

Much of the initial effort within the program has been directed to the development of a core facility and expertise to support a variety of application projects. The program is implemented in the Information Technology Section of the Laboratory where the original staff came from a background in computer-assisted learning and database technology. While there had been some peripheral involvement in AI in the early days of the computer-assisted learning project, this was minimal and there was the immediate need to develop the specialized expertise required to support the new AI program. This was addressed through retraining of existing staff, additions to staff, and through specialized short courses and study visits. As a result, a basic core expertise in applied AI has been developed that now supports a number of application projects.

A second aspect of the development of the core facility is the acquisition of appropriate hardware and software for the program activities. Currently, the available facilities include 11 Sun Workstations, a Symbolics 3620, two MAC-IIs, a MAC-SE, and several PC-type microcomputers (XT, AT, and 386). The Sun workstations, which are the primary hardware resource, are connected in an ethernet along with the LISP machine. The main development tools are ART (Automated Reasoning Tool), Common LISP, Prolog, and C. Other software development facilities include Goldworks, Nexpert, Poplog, C++, Eiffel, Smalltalk, and a variety of small expert system shells. Note that each of these tools represents a significant investment in training and practice for the tool to be effectively available for an application.

Development of the supporting facility is being driven by demands of current and potential application projects, which are identified in consultation with interested users of the technology. Priority is given to those applications which can eventually lead to product development or to improvement of an industrial process. Specific projects may address applications directly or they may have an immediate goal of developing a capability to undertake applications in a new area. In each project, it is desirable to have participation from industry and also from university. The industrial partner represents a client for the technology as well as a source of expertise in the application domain while the university partner brings complementary expertise in the AI field, in the application domain, or in both.

Project selection criteria include the significance of the application, the opportunity represented in terms of access to outside resources, and the potential for inflow of technology, client requirements, and whether or not they can be met with available resources, and the skills and interests of staff. In the remainder of this paper, we describe some of our current activities to illustrate application areas and participation by industry in the program.

Natural Language Interface

The interface to an interactive computer system is determined by the user's requirements as well as the characteristics of the application domain. A common method for soliciting input from the user is a menu interface. However, in applications such as testing and computer-based training this can be unacceptable if the menu provides the user with the clues necessary to formulate his input. In such applications it is highly desirable to have a natural language interface whereby the user can be required to construct his input without uncontrolled aid from the system.

Such an interface is under development for use with a tutorial system shell called TRILOGY [1] produced by Métavision Inc. of Montreal. The natural language interface functions in French and provides the user with flexible communication with the explanation module of the tutor. The goal is to develop a prototype interface suitable for use with a specific subject domain and to specify the methodology required to adapt it to other domains.

Machine Learning

The major element in building an expert system is the task of assembling and coding the required knowledge. Even if an expert is readily accessible, it is often difficult for him to articulate his knowledge in a manner that completely describes his reasoning process. Thus considerable effort has been directed towards automating knowledge acquisition through techniques of machine learning. One technique that has been demonstrated to be effective in practical applications is that of learning by example [2]. Examples can be generated by the expert or more commonly can be selected from a database of past solved cases and used to generate a set of rules that can then be applied to new cases. In addition to generating the initial knowledge base, machine learning techniques can be applied in an incremental mode to adjust the knowledge base according to performance and in this way provide for the system to learn from experience.

Within the NRC program, techniques of machine learning are being investigated in relation to crime analysis and supervisory control of manufacturing systems. In the former, the task is primarily one of identifying patterns from data and developing a structure which can be used to support the identification and diagnostic process. This work was started during a study visit to the Turing Institute where a program was investigated that uses a logic-based approach to simplify the description of the dataset [3]. A second technique [4] from Columbia University involving a generalization-based memory approach is also being investigated. Both techniques are being applied to data supplied by the Ottawa Police on residential break-and-enter incidents. Initial results have demonstrated that meaningful structures can be identified by these methods. The current focus is on improved structuring of the attributes used to describe the events in the dataset.

Application of machine learning to manufacturing systems is being investigated through the development of a manufacturing supervisory system. IMAFO (Intelligent MAnufacturing FOreman) incorporates a mechanism that learns from unsuccessful plans and generates advice to assist the process planner in making performance improvement decisions. The prototype implementation currently uses a process simulator to process commands and generate execution results. The next step is to test IMAFO in a real process planning environment.

Machinery Health Monitoring and Diagnosis

The objective of this activity is to investigate the application of knowledge-based systems to the task of engine health monitoring and fault diagnosis. The project is being undertaken in collaboration with the Division of Mechanical Engineering's Engine Laboratory and is part of a larger study of engine condition monitoring that is being carried out by that laboratory for the Department of National Defence.

The condition of engines and other complex machinery is reflected in parameters such as general performance, temperature, pressure, power output, noise, physical condition of components, and fuel and lubricant consumption. Through new sensor and processing technology it is also possible to consider measures such as oil debris analysis, spectral analysis of exhaust products, and vibration analysis. Information to aid in the interpretation of the performance measures is found in the equipment log, records of scheduled maintenance and overhaul, component replacement history, maintenance manuals, and the experience of the maintenance technician. The total amount and complexity of the information available is such that only a limited portion of it is used by even the experienced maintenance technician to assess the condition of a machine or diagnosing a fault. Artificial Intelligence technology offers a means of combining the available knowledge with data from a particular machine to allow the technician to apply this information more effectively in the assessment of equipment condition. It can guide the technician in the selection of appropriate test procedures, generate a situation assessment based on current test data as well as historical data, and provide for continuity for those tasks that overlap shift boundaries and involve more than one technician. AI technology can also be applied to analysis of historical and test-bed data to discover relationships that can be used in identifying or predicting equipment failure.

The first phase of the project has focused on the development of a prototype to demonstrate feasibility of a knowledge-based system to advise the novice to intermediate engine maintenance technician on trouble shooting an engine fault, identifying the engine subsystem causing the fault, and recommending the required corrective action. The initial knowledge acquisition process uses information from the manuals and from expert maintenance technicians. A preliminary prototype diagnostic shell has been implemented and demonstrated in an application to acceleration problems for the J85 jet engine used in the Canadian Forces CF-5 aircraft [5]. Work has now started towards automating the knowledge acquisition process and extending the project to address the more general problem of condition monitoring.

Realtime Expert Systems

Most expert systems deal with static problems. They are presented with the parameters of a situation and produce an assessment. Realtime expert systems, however, must deal with a continuously changing situation and input values that might be valid for only a limited time. The correctness of the output thus depends on the result of the computation and the time at which the result is produced. But timeliness of the response is not the only factor. The realtime system must be capable of identifying critical decisions and adjusting priorities to suit the immediate situation. The system must be capable of operating continuously and responding to events that do not adhere to any schedule. It must manage its available resources of computing power, time, and memory space. The realtime expert system must incorporate in its knowledge-base information on the structure of the process, the standard operating procedures, the physical design of the plant, expected behavior, remedial actions for known faults, and the heuristics used by operators and system designers. In addition, in the event of a crash, the system must maintain a record of data, decisions, and results of system operations prior to the crash and the ability to use this information for recovery and analysis. The realtime expert system must achieve high performance standards if it is to be applied effectively in areas such as manufacturing and resource industries where potential economic benefits are high.

The NRC is collaborating with MPB Technologies Inc., the Noranda Technology Centre, the University of New Brunswick, the University of Waterloo, McMaster University, Carleton University, Laval University, and the University of Ottawa in the development of an advanced realtime expert system within the framework of the PRECARN initiative. Initial focus is on process control applications but the goal is the creation of a realtime expert system development environment that will be sufficiently general to permit applications in other fields. The first phase of the project consisted of a feasibility study, completed in 1989. The report of this study defined a five-year development project that is one of four projects that have been accepted by PRECARN.

Intelligent Advisor Systems

Systems with embedded computing and information storage capabilities provide a natural environment for the application of knowledge-based systems technology in support of the user. An "Intelligent Advisor," serving as an interface between the user and the system, offers the potential of introducing the new user to the system, providing assistance to the intermediate and occasional user, as well as providing help to the experienced user in special situations. In response to this opportunity, NRC is collaborating with

Laval University, the University of Leeds, Price Waterhouse Quebec Information Systems Group, and Systèmes XA Inc. in the development of an intelligent advisor system for general applications [6]. The objective is to identify and develop basic software that can support the creation and implementation of advisor systems and to develop methodologies for the application of these systems to specific domains and user requirements. The project has taken a bottom-up approach in that a specific topic has been identified for demonstration of feasibility and the solution to this specific problem will subsequently be generalized to a wider set of application domains. The application domain selected for proof of concept is the conceptual design phase of database design methodology. The demonstration prototype of the advisor will be applied to a graphics design aid called Silverrun-ERM. This product of Systèmes XA Inc. supports conceptual database design based on the entity-relationship model.

The Advisor is to support three main functions. First, to introduce neophyte users to the system, it must operate as an intelligent tutor, controlling a structured interaction with the user as defined in a domain-specific knowledge base. Second, it must provide on-line help in response to user requests. This help should assist the user in achieving the immediate goal. Third, the Advisor must monitor a user's interaction with the system, try to identify user goals, and offer, as appropriate, advice based on the user's goals and common design strategies as defined in the domain knowledge base. In each of these modes of operation, the Advisor must transfer the domain knowledge to the user in a form appropriate to the user's skill level and to his immediate requirements.

Support of External Activities

Through the NRC's Industrial Research Assistance Program (IRAP) and the Canadian Space Agency's Space Station Strategic Technologies for Enhanced Automation and Robotics Program (STEAR), the Applied AI Program provides technical guidance to a variety of projects undertaken by Canadian industry. These include the development of an expert system for manpower scheduling by BCW Systems Ltd. of Richmond Hill, Ontario; an expert system medical educator by Knowledge House Publishing Ltd. of Halifax, Nova Scotia; an expert system data analysis tool by Knowledgeworks Research Ltd., Ottawa; knowledge-based learning system development tools and smart simulation trainers by Softwords Division of Press Porcepic, Victoria, British Columbia; space station operations planning by Dynacon Enterprises Ltd., Peterborough, Ontario; and an automated data management system for the space station by Spectrum Engineering Corporation Limited, Kingston, Ontario. Many of these projects are still in the research and development phase but some have gone on to produce marketable products and services based on the work to which the NRC program has contributed. Knowledgeworks is now marketing a software package, KnowledgeSeeker, that enables users to extract decision-making information from a database in the form of a decision tree or as a knowledge base of rules. Knowledge House Publishing Ltd. is marketing an educational package that allows medical students to integrate and apply their preclinical and clinical medical knowledge by managing a simulated patient. Softwords has become one of the most technically advanced commercial sources in North America of knowledge-based learning systems and has supplied systems based on their KBL and Smart Simulator Trainer technologies to Canadian and American government agencies and industrial organisations for training in a wide variety of topics.

Other AI Applications at the National Research Council

The work in the Information Technology Section is focused on the development of generic AI technologies and methodologies. It represents more than 60% of the AI activity currently underway at NRC. Other NRC activities focus on the end application and involve the use of AI technology as a tool to satisfy the requirements of the particular application. In the Laboratory for Intelligent Systems, AI is being applied in the Computing Technology Section in robotics path planning and autonomous vehicle development. In the Robotics and Automation Section in Winnipeg, expert system technology has been applied to generation of procedures for resistance welding, clinical diagnosis of deep vein thrombosis, and assessment of tonsillectomy/adenoidectomy patients for bleeding risk [7]. In the Laboratory for Electromagnetic and Power Engineering, an expert system is being developed for the electromagnetic hardening domain [8]. In the Division of Mechanical Engineering's Systems Laboratory, in Ottawa and

Vancouver, a combination of object-oriented design, database management, and expert systems technologies are being used in a variety of decision support systems. Applications include crew scheduling for a shipping fleet, woods production planning, sky-train simulation, constraint-based scheduling, and a maintenance advisor for a paper making machine control system [9]. The Institute for Research in Construction has established the Advanced Construction Technology Laboratory in Calgary whose mission is to facilitate the introduction of information technology within the construction industry. Specific topics to be addressed by this group include: advanced training techniques, planning systems for implementation and strategic planning, automated design and verification systems, and systems to enable both advanced and novice users access to complex technology. Within the Industrial Materials Research Institute in Montreal, AI techniques are being applied in the design of advanced composite materials. In each of these activities, the AI technology is one of several tools that are being applied to a particular problem. In a few instances, the projects draw on the generic program for support and it is expected that this interaction will grow as the generic program progresses.

Summary

Artificial Intelligence is a pervasive technology which is starting to find its way into a broad range of applications and will have a profound impact on the economy during the next two decades. The NRC's Division of Electrical Engineering began its program in Applied AI in early 1987. The activities of this program are directed towards the development of knowledge-based system tools and methodologies and their application in response to specific requirements. Project opportunities are identified from internal proposals and through consultation with potential users and collaborators within and external to NRC. Priority is given to applications which can eventually lead to product development or improving the effectiveness of an industrial process. Ideally such projects involve an industrial partner, who is interested in and capable of deriving industrial benefits from the project, and academic links that provide a source of expertise. The activity should be in the general area of precompetitive R&D so that the results of the NRC participation can have benefits beyond the immediate project. Specific projects address applications directly or they may have the immediate goal of developing a capability of undertaking future application projects in a given area. In all projects, the industrial link is particularly important and we welcome inquiries from companies regarding joint research and development activities that involve the application of AI technology to real world problems.

References

1. Heath, J., F. Elie, and M. Habib. The educational potential of rule-based expert systems in the micro environment. *Proceedings of the Sixth Canadian Symposium on Instructional Technology*, Halifax, Nova Scotia, 1989. pp. 85–89.

2. Quinlan, J.R., *et al.* Inductive knowledge acquisition: a case study. In *Applications of Expert Systems*. J.R. Quinlan (ed.). Addison-Wesley, Sydney, 1987.

3. Muggleton, S. Duce, an Oracle based approach to constructive induction. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, 1987. pp. 287–292.

4. Lebowitz, M. Experiments with incremental concept formation: UNIMEM. *Machine Learning*, 2, 103–138 (1987).

5. Abu-Hakima, S., *et al.* Jet engine technical advisor (JETA). *Proceedings of the Second International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, Tullahoma, TN. June 6–9, 1989. pp. 154–160.

6. Boulet, M.-M., *et al.* A design task advisor. *Artificial Intelligence and Education, Proceedings of the 4th International Conference on AI and Education*, Amsterdam, May 24–26, 1989. pp. 25–31.

7. Pizzi, N.J., S. Kapoor, and J.M. Gerrard. Hematology expert system (HES) for tonsillectomy/adenoidectomy patients. *Proceedings of SPIE Applications of Artificial Intelligence VII*, Orlando, FL. March 27–31, 1989. pp. 884–893.

8. LoVetri, J., *et al.* HardSys: Applying expert system techniques to electromagnetic hardening. *Proceedings of the IEEE Symposium on*

Electromagnetic Compatibility, Denver, CO. May 23–25, 1989. pp. 102–115.

9. Evans, S., S. Nuyan, and R. Amyot. Expert system diagnosis in CD weight control. *Proceedings of the 76th Annual Meeting of the Technical Section, Canadian Pulp and Paper Association*, Montreal Quebec. Jan. 24 – Feb. 2, 1990. pp. A231–A235.

For further information, please contact:

J.W. Brahan
Institute for Information Technology
National Research Council
Ottawa, Ontario
K1A 0R6
Telephone: (613) 993-2484
Fax: (613) 952-7998
jack@ai.iit.nrc.ca

Provincial Research Organizations

By Ernest Chang, Alberta Research Council

Introduction

In writing this report, each member of the Association of Provincial Research Organizations (APRO) was contacted. Only those PROs which reported activity, either verbally or in writing, are included below. As much as possible, the report for each PRO is taken verbatim from written material. It should be noted that these reports represent the activity only of the members of the APRO and not all government-based activities within the provinces. For example, although the British Columbia Research Council did not report any AI activity, the British Columbia Advanced Systems Institute does support AI activity through sponsorship of university and industrial organizations.

New Brunswick Research and Productivity Council

At the New Brunswick Research and Productivity Council researchers are using fuzzy logic to implement some portions of a control/alarm/advisory system for drag lines used in open-pit mining. They have also used fuzzy logic in a microprocessor-based temperature control system.

For more information, please contact:

Lee A. Danisch, P. Eng.
Department Head, Electronics
Research and Productivity Council
P.O. Box 20 000
Fredericton, New Brunswick
E3B 6B2
Telephone: (506) 452-8994
Fax: (506) 452-1395

Centre de recherche industrielle du Québec (CRIQ)

This section describes only work at the Montreal Laboratory.

Activities in Artificial Intelligence

Machine Vision

Current activities in the field of machine vision are in the areas of automated visual inspection and robotics. A project presently in the proposal stage concerns the application of AVI techniques to small machined metal parts. Image analysis techniques have also been used to identify parasites in fillets of fresh fish.

The machine vision team works with various sets of hardware including Sun stations. A dedicated laboratory is available.

Expert Systems

Expert system activities are presently at a low level due to the return to the academic world of some of our key people. However, we maintain an active interest in the field, in particular with respect to applications in the field of manufacturing automation, the biomedical domain, and intelligent database front ends.

Neural Nets

Activities in this field are very recent, following the acquisition of an HNC neural net simulator. We are using the neural net techniques in one industrial application for character recognition.

For more information, please contact:

Carole Rach
Centre de recherche industrielle du Québec
8475, rue Christophe-Colomb
B.P. 2000, Succursale Youville
Montreal, Quebec
H2P 2X1
Telephone: (514) 383-1550
Fax: (514) 383-3250

Saskatchewan Research Council

The following is from a telephone conversation with Ravi Mathiel of the SRC.

There are a few projects identified in the area of AI applications, but none of them have started. There are no commitments at the moment, but the technology is being tracked, and prices are coming down. Projects will definitely be applications oriented.

For more information, please contact:

Ravi Mathiel
Saskatchewan Research Council
15 Innovation Boulevard
Saskatoon, Saskatchewan
S7N 2X8
Telephone: (306) 933-5499
Fax: (306) 933-7446

Alberta Research Council

Introduction

The Advanced Technologies Department was formed in 1984 with a specific mandate to apply techniques from artificial intelligence, robotics, and automation to problems originating in the industrial sector. As a technology transfer agency, the Department has worked with private sector and other organizations primarily through the contract research route.

In 1985, the Department had about 10 technical staff; today, the Department has more than 40 employees and has been renamed the Advanced Computing and Engineering Department. In 1989, the Department executed about 25 contracts with revenues of $1.8 million. Most of these were in applications of AI, robotics, and automation, though the Department also worked in other areas of advanced computing, such as computer-supported collaborative work.

Through the existence of the Department, it has conducted over 40 projects in AI, robotics, and automation. Most of its clients are within Alberta, but it has also worked or teamed with companies from British Columbia, Saskatchewan, Manitoba, and Ontario. It is one of the principle suppliers of technology to the Strategic Technologies for Enhanced Automation and Robotics (STEAR) program of the Canadian Space Agency, and is a member of PRECARN. It has worked with companies and organizations such as Hewlett Packard Laboratories, Northern Telecom, Alberta Government Telephones, Environment Canada, SED Systems Inc., MacDonald Dettwiler and Associates Ltd., IBM Don Mills, and the Department of National Defence of Canada.

The Department also conducts an outreach program as part of technology dissemination, conducting seminars and workshops in expert systems, robotics, and advanced manufacturing. It has close collaborative ties with the AI activities of the National Research Council and is working jointly in the area of AI in construction with the Advanced Construction Technologies Laboratory (Calgary) of the Institute for Research in Construction of the NRC.

Canadian Artificial Intelligence, a quarterly magazine, is an official publication of CSCSI, the Canadian Society for Computational Studies of Intelligence. It has been edited and produced by staff of the Advanced Computing and Engineering department and the Alberta Research Council for the past two years.

Technology Groups

The Advanced Computing and Engineering Department is organized into five technology groups.

Artificial Intelligence in Operations

The Artificial Intelligence in Operations group promotes the identification, acquisition, development, and dissemination of artificial intelligence technologies applied to operations problems. The group uses these technologies: expert systems, constraint-based reasoning, temporal reasoning, qualitative modeling, operations research, and simulation. Applications for this area of work occur in a number of sectors: manufacturing; design, planning, scheduling; transportation; fleet scheduling; aerospace; fault diagnosis, maintenance; environment; and impact assessment, accident management. Other applications are process planning, modeling, and design.

The group plans to invest resources in a number of internal R&D initiatives: investigation of intelligent design development of a software library of constraint satisfaction algorithms, implementing a program of technical awareness of appropriate traditional techniques, and monitoring of developments in AI.

The group comprises nine technical staff (including the group leader). Six members have a M.Sc. in AI, one has a Ph.D. in Manufacturing Engineering, and two have a B.Sc. (Engineering and Computing). All have experience on research contracts and joint ventures with industry.

Learning and Collaboration

The Learning and Collaboration group joins three complementary fields: collaboration technologies, intelligent learning systems, and collaborative learning systems. Collaboration technologies address the computer's potential to enhance the group work process through direct support of communication needs and specific collaboration requirements across space–time domains.

Intelligent learning systems focus on creating intelligent interactive learning environments that manage the learning process and deliver mediated instruction. They include intelligent tutoring systems (ITS), computer-guided discovery environments, interactive simulation, and computer-managed learning (CML) systems. The fusion of collaboration technologies and intelligent learning systems extends computer-mediated learning systems to include collaborative learning, in which teams of persons can learn coordinated tasks.

This group has considerable academic expertise and technical experience. Group members have three Ph.D. (computer science), three M.Sc. (special education, computer graphics, engineering), three B.Sc. (computer science, psychology). Three of the group members are adjunct professors at the Universities of Calgary and Saskatchewan.

Pattern Analysis and Machine Intelligence

The Pattern Analysis and Machine Intelligence group focuses on data interpretation and the emerging field of machine learning. Activity areas include sensors and vision, syntactic pattern recognition, and connectionist systems, as well as intelligent systems in machine learning. Techniques include knowledge representation, inference mechanisms, genetic algorithms, case-based reasoning, and machine discovery.

Future plans include employing machine-assisted learning techniques, providing value-added interpretation of data from imaging satellites, as well as developing expert systems with the inclusion of computer-based knowledge acquisition tools.

Group members' educational qualifications consist of four M.Sc. (computer vision, computer science) and three B.Sc. (computer science, mechanical engineering).

Realtime Artificial Intelligence

The Realtime Artificial Intelligence group focuses on applying AI methodologies and advanced programming techniques to enhance the capabilities of realtime systems for monitoring and control. These methodologies and techniques include constraint-based reasoning, qualitative reasoning, truth maintenance, knowledge representation, temporal reasoning, object-oriented programming, structured design methods, intelligent simulation, and modeling. The principal benefit of using realtime knowledge-based systems (KBS) is to reduce the cognitive load on users or to enable increasingly complex systems to be effectively managed.

Capabilities which will be enhanced by the use of these include alarm handling, trend analysis, fault diagnosis, and dynamic data acquisition and processing. The sectors which apply these technologies include oil and gas, utilities (pipeline, electric, telecommunications), defence, health, and aerospace.

Group members have current experience developed through existing projects and research. Two have M.Sc. (computer science, engineering) and two have M. Math.

Robotics and Automation Technologies

The Robotics and Automation Technologies group aims to integrate advanced computing with real mechanical systems. This includes design and implementation of manufacturing systems and work cells, control system, planning systems, and robotic manipulators. The group works in these technologies: robotics, automation, sensor fusion, process control, and signal processing.

Application areas within the experience of the group include: control systems design, development, and implementation, from applications of PLCs, through controller design, up to hierarchical and intelligent systems: robotic work-cell design and development including 3-D simulation: sensor integration for robotics applications (proximity, tactile, force, torque, vision): computer-integrated manufacturing (CIM), CNC machining, material/parts transportation and handling, and assembly operations. Past work has been in these sectors: manufacturing, health care, space, natural resources, agriculture, and the entertainment industry.

The group consists of a Ph.D. in Electrical Engineering (digital signal processing), an M.Sc. in Mechanical Engineering (robotics and dynamics), an M.Sc. in Computer Science, a B.Sc. in Mechanical Engineering, and a B.Sc. in Computer Science.

Resources

These include a DEC MicroVAX electronic mail server, a network of Macintoshes, IBM PCs, Sun, and Symbolics workstations. Tools for developing AI applications and expert systems include KnowledgeCraft, Automated Reasoning Tool (ART), and Knowledge Engineering Environment (KEE), as well as Nexpert Object, Goldworks, and Quintus Prolog. Computer-aided design and manufacturing (CAD/CAM) software includes AutoCAD, Silma Cimstation (robotic kinematics simulation), and Smartcam.

The Department also operates an advanced manufacturing demonstration facility that allows investigation of CAD/CAM, robotics, sensor fusion, and integrated computer control, incorporating a Puma and an Adept robot, a conveyor system with customized fixtures and a 4-axis CNC milling machine.

Projects

The following are selected examples of recently completed projects involving AI:

- *Alberta Government Telephones*: This was a contract with Alberta Government Telephones (AGT) to develop a configuration expert system on a PC for a small telephone switch to be marketed from retail phone centres.
- *Alberta Transportation*: This was a contract with Alberta Transportation to create an expert system to assist in designing intersections for secondary highways. The system has been completed and is presently being integrated into the client site.
- *Biotechnology Department*: This is a joint project with the Biotechnology Department of the Alberta Research Council to develop and implement a prototype scheduling system for automated scheduling of their fermentation plant, using constraint-based techniques.
- *Computer-Based Training Systems Ltd.*: This was a joint research venture with Computer-Based Training Systems Ltd. (CBTS), Calgary, to develop an expert computer-managed learning (CML) system that includes curriculum design and course development, and an exam-generation component.
- *D&S Knowledge Systems Inc.*: This was a joint research venture with D&S Knowledge Systems Inc., Calgary, to develop a well log analysis expert system. The project integrated an expert system with LOGMATE, a fourth-generation log analysis workstation. Together they provide graphics, database, data communications, arithmetic, and report generation facilities. Given well-log data, the expert system will either automatically provide an analysis of the data or will suggest data specific analysis methods for the user to invoke. The system is implemented in C and runs on an AT compatible and on Unix-based microcomputers.
- *Dataspan Technologies Inc.*: This is a joint research venture with Dataspan Technologies Inc., Calgary, to automatically interpret handwritten digits and symbols on maps and map overlays. This knowledge-based system for character/word recognition augments existing raster-vector conversion software and vector editing facilities. The digit classification system was enhanced by integration of an AI component that uses local contextual knowledge, the known format for the representation of numbers, and expectations of the values of numbers in particular documents to resolve digit ambiguities and to concatenate adjacent digits to form numbers. The final system is being implemented in C on a Sun computer. The success of this project will be significant to the artificial intelligence community and to the document conversion industry. The project has also led the company to change its operations — convert to an advanced computer language and incorporate a software revision control system and project management tools to increase their productivity.
- *DND/CF-18*: This is a contract with the Department of National Defence to develop an expert system to assist with the maintenance and fault diagnosis of the CF-18 aircraft. The contract started in April 1990.
- *Det norske Veritas/Centre for Frontier Engineering Research*: This is a subcontract with Det norske Veritas, now transferred to the Centre for Frontier Engineering Research, Edmonton, to develop an intelligent monitoring system for the ballast control system of semi-submersible oil drilling units, in partial fulfilment of a STEAR contract with the Canadian Space Agency.
- *Expert system for well test interpretation*: This was a contract to determine the feasibility of developing an expert system for well test interpretation. A prototype qualitative modeling mechanism for interpreting graphical data, which identifies characteristic curve shapes that correlate to physical phenomena, was designed and implemented.

- *HERMES (Heuristic Emergency Response Management Expert System)*: This is an expert system that provides advice on the management of spills of hazardous materials. The feasibility prototype was constructed with Alberta Public Safety Services and Emergency Preparedness Canada as clients. The follow-on contract is funded by a consortium of five provincial and federal government departments, and will extend the capabilities of the original prototype.

- *International Submarine Engineering Ltd.: Power Management*: This is a subcontract with International Submarine Engineering Ltd., Port Moody, British Columbia, which was awarded a contract by the Canadian Space Agency's STEAR program. The contract is to develop the conceptual design for the power management system for the Mobile Servicing System (MSS) on the proposed Space Station. Artificial Intelligence is used to assist in the fault diagnosis component.

- *ISIM, an Intelligent Simulation System*: ISIM was designed to integrate reasoning and representation techniques developed in artificial intelligence with discrete-event simulation. A simulation environment was created by enhancing KnowledgeCraft with a Simula/DEMOS-like representation methodology. The ISIM specification environment is augmented by a number of tools: traditional quantitative analysis techniques such as response surface methodology (RSM), and qualitative dependency analysis. The expert system may use these tools to run an optimal set of experiments for analysis of the specified model.

- *Northern Telecom*: This is a series of contracts with Northern Telecom, Toronto, to implement an expert system for process planning of printed circuit board component insertions. The expert system is now being deployed to other plants across Canada.

- *PharmHand*: This was a project with Red Deer Regional Hospital Centre to develop a robotic workcell for reconstituting intravenous medication stored in powdered form. The feasibility prototype has been constructed and the hospital is seeking additional funding for developing a production prototype.

- *Process Planner π*: Process Planner π is a tool that can be used for process planning, and is based on the specification, planning and execution of changes in physical systems involving time. The system draws on research from temporal reasoning, planning, knowledge representation, and natural language understanding. The system has two main parts, a representation of the environment and a planner.

- *SWIFT, an expert system for severe weather forecasting*: This was a contact with MacDonald Dettwiler and Associates Ltd. and the Canadian Atmospheric Environment Service. The system assists with short-term severe weather forecasting and the accurate issuance of severe weather warnings. SWIFT permits integration of conventional systems including databases, with a new weather representation system. One of the modules, METEOR, is a severe-weather expert system previously developed at the Alberta Research Council by Renee Elio and John DeHaan. This rule-based expert system accepts information about meteorology and geography, and produces an assessment of the situation and some advice. SWIFT demonstrates the integration of large, frequently updated databases with the reasoning capabilities of an expert system. It was implemented on a Xerox computer, in InterLISP, LOOPS, and OPS4.

For more information on the Advanced Computing and Engineering department and the Alberta Research Council, please contact:

Dr. Ernest Chang, Department Head
Chris Lumb, Manager, Industrial Liaison
Advanced Computing and Engineering Department
Alberta Research Council
3rd Floor, 6815-8 Street N.E.
Calgary, Alberta
T2E 7H7
Telephone: (403) 297-2600
Fax: (403) 275-3003
INTERNET: echang@noah.arc.ab.ca
lumb@noah.arc.ab.ca

Toll Free Technology Information Line
1-800-661-2000 (in Edmonton: 450-5000)

Transport Canada

Purpose

To review Transport Canada's experience in Artificial Intelligence Research between 1983 and 1990 and to outline future directions.

Introduction

The transportation sector is technology intensive. In this context Transport Canada's operational responsibilities, for example for air systems navigation, and its regulatory responsibilities for safety and security, require the department to assess the feasibility of applying new technologies and, where appropriate, to stimulate or otherwise or support their application. The department maintains an in-house core of technical expertise for this purpose. The activities of this core are supplemented through contacts and cooperative research with other federal agencies, notably the National Research Council, and other Canadian and foreign transportation technologists. The department contracts out over 95% of its R&D program. This amounts to between $25 and $35 million annually.

The department's R&D program priorities are designed to support Transport Canada's program priorities. R&D for the Operational Groups includes, for example, human factors research in support of Air Traffic Controller workstation design, de-icing R&D, oil spill cleanup R&D, ice breaker design, and, R&D in support of the department's large ($500 million+ annually) capital investment plan. The department also supports R&D designed to increase safety in the sector, for example, Advanced Train Control System R&D with the railways, Intelligent Vehicle Highway Systems R&D with the provinces and American government agencies. Underpinning these activities is a commitment of resources to monitoring technological developments, and, where appropriate, to promoting their application through "Exploratory" R&D.

The department's initial artificial intelligence (AI) R&D activities commenced through the exploratory technology monitoring process in 1983. The department's AI R&D continues to be exploratory in areas such as neural networks. Conversely, through technology transfer, diffusion, and an explicit programs to establish a capacity to "manage" applications, the development and implementation of relatively complex expert systems is now widespread throughout the department and becoming more so in the sector. The remainder of this paper demonstrates how and why this occurred.

Experience 1983–1986

From the late seventies on, faced with the incredible explosion of information technologies, the department assigned high priority to R&D in support of advanced computer, communication, and control systems. (This imperative was strengthened by the inability of conventional transportation technological solutions, e.g., faster planes and trains, to cope by themselves with emergent transportation issues.) By 1983 it had become intuitively obvious that if expert systems could provide accurate medical diagnoses, similar diagnostic systems could and should be developed for transportation. The development of the G.E. Delta Cats Expert System for rail engine maintenance, and some preliminary work on highway and transportation of dangerous goods expert systems, supported this intuition.

The department's initial research culminated in a seminar, held in Montreal in September 1985, designed to bring AI specialists and transportation professionals together to discuss preliminary analyses of opportunities for applying AI to transportation, and to advise the department on next steps. (The report of this seminar is listed in the bibliography at the end of this note.)

A paper by Paul Tuan, Stanford Research Institute, on an expert system for enhancing vehicle dispatch activity and an "Expert System Primer" by T. Gomi of Applied A.I. and R. Sylvester of Cognos, provided the basis for the seminar discussion. Perhaps the most important contribution of this seminar was to replace scepticism and disbelief by enthusiasm. It also established the principles which have guided the evolution of the department's AI R&D since that time:

- AI systems support decision making. Expert systems are a "skilling" technology.
- Virtually every transportation operational and regulatory task can be effectively and economically supported by an expert system. There are important niche markets for neural net,

voice recognition, natural language processing, intelligent tutors, and other variants of AI in transportation.

- Expert systems are the most mature of all AI technologies. Notwithstanding, it would take at least two years (from 1985) to build up a capability to manage expert systems R&D and to develop and demonstrate "simple" expert system applications.
- It is important to "begin." The process of managing expert system projects is an essential teaching mechanism.

The publication of the report, *Expert Systems: Their Application in the Canadian Transportation Sector*, in January 1986 effectively completed the department's first steps in AI R&D. From now on the emphasis would be on expert systems development, on monitoring developments in other fields of AI, and, where and when required, developing applications of other AI technologies.

Experience 1987–1990

The potential areas for application of expert systems to transportation were established as:

- Vehicle Crews
- Vehicle Maintenance
- External Control of Vehicles in transit
- Permanent Way control and operations
- Terminal control and operations
- Operational Planning and Regulation
- Strategic Planning

Each area was further subdivided to permit identification of specific, "do-able," useful projects. The "Vehicle Crews" category, for example, was subdivided into: vehicle control, vehicle system diagnostics and trouble shooting, on-board navigation, communications, cargo monitoring and service, emergency advice, specialized, e.g., fire-fighting on-board applications, and procedural/regulatory advice. The Permanent Way Control and Operations category consists of scheduling and capacity planning, communication and control, facility diagnostics and maintenance, weather forecasting, usage accounting, emergency procedures, and procedural/regulatory advice.

In 1986–1988 a number of relatively simple expert systems were developed. Paul Tuan's dihspatch/scheduling product was applied to a para-transit operation in Vancouver and then in Edmonton. A project to assist the air traffic controller collision avoidance function for "sparse" airspace was completed. Marine inspection activities in the Coast Guard were supported through a small but effective expert system. The projects were useful, they developed competence, and their benefits could readily be demonstrated to management.

By 1988 it had also become evident that the results of Transport Canada's AI R&D program had exceeded all expectations. (Some 15 expert systems had been developed or were approaching completion. It had been expected that five would have been completed.) A second seminar was therefore held. It was designed to confirm lessons learned, and to establish guidelines for Transport Canada's AI R&D to the mid-nineties. The number of technical papers produced for this seminar confirmed the earlier belief of the broad applicability of AI to transportation. The list included:

- An Expert System for the Resolution of Interference between Radio Systems
- HERMES: An Expert System for the Management of Spills of Hazardous Materials
- A Rule-based Highway Location System
- Signal Abstraction for Machine Analysis of Radar Signal Profiles
- Sixth Generation Computing: Natural Language Interfaces, their Limitations and Future Prospects
- Neural Nets: Beyond Knowledge-Based Systems
- An Expert System for Air Traffic Control Training
- SWIFT: An Expert System for Severe Weather Forecasting
- Future Trends in Road Traffic Control

More details on these papers and the six workshops which were held on the second day of this seminar are contained in the *Proceedings*. The technical results confirmed the directions established for the program in 1986. The discussion also revealed that concerns about the ability of the Canadian AI community to meet transportation demands for technical support were unfounded.

Currently Transport Canada is spending over $2 million annually on AI R&D, primarily for expert system development. In addition the Aviation Group has established an AI Testing;AItesting and evaluation facility for the department. Exploratory R&D in neural nets, voice recognition, and intelligent tutors suggests that these technologies are approaching the point where they can be applied in transportation.

Conclusion

The scope and diversity of AI research by Transport Canada is such that its technical dimensions, the opportunities and issues, cannot be captured in one paper. The paper has thus provided the reader with an indication of how AI research is viewed in Transport Canada, and transportation, where it has been and where it is going.

Bibliography

Proceedings of the Workshop on the Application of Expert Systems to Transportation. September 12, 1985. TP 7209E.

Expert Systems: Their Application in the Canadian Transportation Sector. January, 1986. TP 7328E.

AI-II. Proceedings of the Second Workshop on the Application of Artificial Intelligence to Transportation. March 28–29, 1988. TP 9444E.

For further information, contact:

Mr. R.T. Lewis
Research and Development
Transport Canada
Tower C, 29th Floor, Place de Ville
330 Sparks Street
Ottawa, Ontario
K1A 0N5

University Research in AI: An Overview

Introduction

Long before the artificial intelligence boom of the eighties, universities were incubating the discipline in Canada. Among the universities with long-standing interest in AI (dating back to the sixties) are the University of British Columbia, The University of Western Ontario, the University of Alberta, and the University of Toronto. Since the mid-seventies there has been a steady growth of AI capabilities in universities, until today most Canadian universities have at least one faculty member with research interest in artificial intelligence. Some universities have many professors, graduate students, and research assistants covering a variety of AI subspecialities; others have made a smaller commitment to the discipline having one or two professors with AI as a primary or secondary research interest.

A Brief Idiosyncratic Overview of Canadian Academic AI Capabilities

Canada's academic AI capabilities are second to none. All areas of artificial intelligence are represented. Most of this research is centered in computer science departments, although increasingly AI is spreading its tentacles widely throughout the university. What follows is a very brief overview of the kinds of research underway in Canadian universities; inevitably certain projects and programs will have been missed, but the synopsis is accurate in general terms at least.

Computational vision is one of the main areas of strength, with research ongoing in Canadian universities coast to coast, most notably University of British Columbia (UBC), Simon Fraser University, McGill University, and University of Toronto, but also in many other places like Dalhousie University, University of Saskatchewan, and elsewhere. Applications include satellite photo interpretation, analysis of motion, sketch map interpretation, examination of logs for possible cutting in sawmills, among many others.

Another major area of research interest is knowledge representation, with a superb group at University of Toronto, major concentrations at Simon Fraser University and University of Alberta, and many other researchers from Acadia to Victoria, from Guelph to Saskatchewan, from Waterloo to UBC. Canadian academics are among the world leaders in formal knowledge representation techniques, but there is also work in more applied styles of knowledge representation such as logic programming, semantic networks, frames, and production systems.

Natural language understanding is another one of the main areas of research in Canada with expertise at Simon Fraser University, University of Toronto, University of Waterloo, University of Manitoba, and McGill University among other places. There is interest in parsing, semantics, discourse and dialogue modeling, and speech understanding. Canadian academic interest in cognitive science is also high, with a major center at The University of Western Ontario, and other important locales throughout the country. Learning and knowledge acquisition is also of importance in many institutions in Canada, notably the University of Calgary, but also including the University of Regina, Simon Fraser University, and the University of Toronto, among others. The University of Calgary is also a Canadian center for human–computer interaction, although there is important activity in this realm elsewhere as well. Intelligent tutoring systems research and other investigations into AI and education are widespread in Canadian academe, with special focus at University of Saskatchewan, Université de Montréal, and Laval University, but also in smaller institutions like Athabasca University, Université du Québec à Montréal, Université du Québec à Chicoutimi, Trent University, and Memorial University. Research into AI programming environments is ongoing, most notably at Carleton University and Queen's University.

Robotics is also being investigated at Canadian universities like McGill, usually in various engineering departments rather than computer science departments. Expert systems research is ongoing in universities, sometimes in computer science departments (such as University of Ottawa), but just as often in other departments with particular applications requiring expert systems technology. An interesting approach to helping expert systems technologies disseminate throughout a university has been taken by Simon Fraser University which has an Expert Systems Laboratory set up for this purpose, among others. The ESL is partially privately funded and is coordinated through the Centre for Systems Science which is a consortium of a dozen university schools and departments. The social impact of artificial intelligence has also been a concern of Canadian academics, notably at UBC, Waterloo, and Toronto. A previous report of the Associate Committee on AI outlines some implications of AI for Canadian society .

Notable Canadian Academic Success Stories

CSCSI/SCEIO — The Canadian AI Society

In 1973 there were some 30 AI researchers in Canadian universities and only a few graduate students doing research in the area. There was little or no industrial AI activity, either in Canada or elsewhere in the world. Nevertheless, these hardy pioneers felt that the discipline would grow and that the common interests of Canadian AI would be well served by the formation of a society to coordinate these mutual interests. Thus was born the Canadian Society for Computational Studies of Intelligence/Societe canadienne pour l'étude de l'intelligence par ordinateur (CSCSI/SCEIO), formed in 1973 as the first national AI society in the world (ACM's SIGART and the British-based AISB both already existed but had international mandates).

Within a year the society had held a workshop in Ottawa, the precursor to the biennial conferences which were instituted in 1976 at UBC. The circle has been completed with the Eighth conference being held in Ottawa in 1990. The CSCSI/SCEIO conference is now a major one on the international scene. It has very high standards — indicative of this is that the best paper from the conference is invited to appear in the *Canadian Artificial Intelligence* journal, the oldest and one of the pre-eminent international journals in the field. The conference is also traditionally a friendly interactive one, without the vast number of parallel sessions which make it so hard to get anything out of the big international conferences.

Shortly after its inception, the CSCSI/SCEIO also instituted a newsletter, which appeared annually for several years before appearing more regularly after 1979 and being transformed into the current glossy quarterly magazine in the mid-eighties. The CSCSI/SCEIO is the sponsoring society for the *Computational Intelligence* journal, published by the National Research Council of Canada. It has also hosted workshops and has been a sponsoring agency for many national and international conferences. For administrative purposes the society is now affiliated with the Canadian Information Processing Society (CIPS). From its small academic beginnings the society has grown to 700 members from universities, industry, and government, and is well respected internationally as the representative of Canadian AI.

CIAR Project

In the mid-eighties, the Canadian Institute for Advanced Research (CIAR) was founded to promote advanced, long-term research in Canada through the funding of networks of internationally well-known scholars at Canadian universities. Scholars funded by the CIAR stay at their own university but are paid entirely by the Institute for five years and are freed from the majority of their teaching duties so as to concentrate on research. Regular project meetings and electronic networking attempt to foster synergy among the CIAR fellows.

The first project funded under the auspices of the CIAR was the AI, Robotics and Society (AIRS) project, a testament not only to the importance of AI to Canada and the interesting long-term research issues being tackled by AI, but also to the quality of Canadian academic AI scholarship. The AIRS project started with scholars at three nodes: UBC, Toronto, and McGill. It has since expanded to include fellows and associate fellows from other universities. The project is currently half-way through its second five-year mandate. The CIAR project is described in more detail in another section of this ACAI committee report.

The IRIS Project

One of the 14 projects funded under the Canadian government's Networks of Centres of Excellence program is the IRIS project. IRIS consists of a large interconnected group of AI researchers scattered through many of Canada's universities. With major emphases on computational vision, robotics, and knowledge-based systems, the four-year project should infuse some $23.8 million into the research programs of many Canadian academic AI researchers. The networks competition was a fierce one, and several hundred initial applications had to be winnowed to the 14 which had world class status and high potential. It is a tribute to Canadian academic AI excellence that among the 14 was a project whose major emphasis is artificial intelligence.

International Successes

Canadian AI researchers have made a major impact on the international AI scene. Many graduates of Canadian universities are serving in universities, research laboratories, and corporations in the United States and overseas. Canadians have won awards on the international AI scene, including best paper awards at major conferences. Of particular significance is the fact that the Computers and Thought award for the most outstanding young scholar given out biennially at the International Joint Conference on Artificial Intelligence has been won the last three times by people with at least one degree granted by a Canadian university (Hector Levesque with B.Sc., M.Sc. and Ph.D. from University of Toronto; Johann deKleer with a B.Sc. from UBC; and Henry Kautz, an American native with a M.Sc. from University of Toronto). Canadians have high standards of scholarship and, in fact, had the highest acceptance ratio of any national grouping at last year's IJCAI conference in Detroit. Canadian academics serve on many international program committees, appear on the editors' boards of major AI journals, are frequently invited speakers at international conferences and workshops, and often consult at the national and international level.

Some Case Studies

Descriptions of several different university-based AI research endeavours follow this article.

For further information, please contact:

Gord McCalla
ARIES Laboratory
Department of Computational Science
University of Saskatchewan
Saskatoon, Saskatchewan
S7N 0W0

Nick Cercone
Centre for Systems Science
Simon Fraser University
Burnaby, British Columbia
V5A 1S6

Research in the Knowledge Sciences at the University of Calgary[1]

By Ian Witten and Brian Gaines[2]

Since the late 1960's the Computer Science Department at the University of Calgary has enjoyed a major research presence in the field of human–computer interaction, and the last 5 to 10 years have seen sharply increased emphasis on AI topics and methods. This trend was focused and strengthened by the establishment in 1985 of the Knowledge Science Institute (KSI), whose objectives are to study knowledge science, technology, economy, and social impact, and to promote interdisciplinary research on knowledge-based systems. This has involved the provision of space, capital equipment, operating funds, the attachment of faculty, and the recruitment of research staff.

While the KSI operates as a network of multidisciplinary research groups spanning several faculties and academic disciplines, this overview concentrates on knowledge science research undertaken by, or in close collaboration with, computer science faculty members. It is divided into the areas of knowledge acquisition for expert systems, machine learning and programming by example, system theory and organization (especially modeling and prediction), and logic programming. However, different projects mingle and overlap, and in practice no real distinction is made between those that involve artificial intelligence techniques and those that do not. Research in allied fields of computer science, such as human–computer interaction, distributed systems, and formal methods for software and hardware verification, is not covered here.

Research in Knowledge Engineering

A model of the development of the computing industry has been developed and used to account for past events and predict future trends, particularly fifth and sixth generation priorities. The industry came into being as a result of electronic device technology, thus enabling society to cope with an increasingly complex world. Underlying all developments in computing is a tiered succession of learning curves which make up the infrastructure of the computing industry. The model provides a framework for the industry based on this logical progression of developments. It links the framework empirically to key events in the development of computing. It links the framework theoretically to a model of economic, social, scientific, and individual development as related learning processes with a simple phenomenological model. The fifth generation development program, with its emphasis on human–computer interaction and artificial intelligence, and the sixth generation program, with its emphasis on knowledge science, are natural developments in the foci of attention indicated by the model [9, 10]. Knowledge engineering represents a key technology within this model, and is a cornerstone of our activities.

Rapid Prototyping

Knowledge acquisition techniques are being developed for rapid prototyping of knowledge-based systems [14, 15, 16]. Problems of knowledge engineering have been recognized since the early days of expert systems and are impeding practical applications. It is difficult to elicit from experts the knowledge structures underlying their skills. This has heightened interest in the possibility of providing knowledge support systems to automate knowledge engineering as a process of direct interaction between domain experts and the computer [6, 8]. In 1980 we proposed that personal construct psychology could provide foundations for expert systems, particularly in systems that combined interactivity with database access and expert advice to provide decision support, and gave examples of algorithms and programs that extracted entailment rules from repertory grid data. In 1983, we reported further enhancements of these techniques and a preliminary experiment to validate them empirically as a knowledge engineering technique for priming expert systems. This work led to industrial studies of the methodology applied to the development of expert systems: Boeing Computer Services and Lockheed Software Technology Center have reported success in applications; and validation has been reported in a statistics domain.

The knowledge acquisition research focuses on the use of techniques based on personal construct psychology to automate knowledge engineering for expert systems. Knowledge support systems are being developed as interactive knowledge engineering tools, and the design criteria for such systems are being clearly defined. KITTEN is a knowledge support system developed for AI workstations which includes tools for interactive repertory grid elicitation and entailment analysis that have been widely used for rapid prototyping of industrial expert systems. It also includes tools for text analysis, behavioral analysis, and schema analysis, which offer complementary and alternative approaches to knowledge acquisition. The KITTEN implementation integrates these tools around a common database with utilities designed to give multiple perspectives on the knowledge base [17, 18].

---

[1] Reprinted with permission from *Canadian Artificial Intelligence*, Oct.: 25–30; 1988.

[2] Ian Witten and Brian Gaines are professors at the University of Calgary. Brian Gaines is founder and director of the KSI.

A wide range of tools developed elsewhere are also being investigated such as AutoIntelligence, ETS, KSSO, and Nexpert; and validation studies are being undertaken of the performance of these tools in eliciting and structuring expert knowledge.

Hypermedia and Knowledge-based Sports Coaching

In the Knowledge Science and Neuro-Motor Laboratory in the Faculty of Physical Education, studies are being undertaken of knowledge-based sports coaching systems and the cognitive foundations of knowledge transfer [19]. The research focus is on eliciting, understanding, and applying the knowledge structures underlying the performance skills of athletes and the coaching skills of those involved in their selection and training. The abstract knowledge structures for sports coaching have been developed independently of specific sports. They have been shown to lead to significant analyses of a wide range of sports that guide the generation of coaching strategies and the individual acquisition of playing skills, and to be meaningful and comprehensible to both coaches and players. These knowledge structures are being applied through a variety of media, including books and computer-based hypermedia systems [20].

The new *Steps to Success* book series from Human Kinetics Publishers is based on the explicit use of the structures, and is available in handbooks for athletes and coaches. Initially 10 sports are covered, including archery, badminton, basketball, golf, gymnastics, tennis, and swimming. In the research laboratory a range of hypermedia systems based on the knowledge structures are being developed, including an expert system for ice hockey coaching which was shown at the Winter Olympics in February 1988. This runs on a Macintosh using specialist knowledge acquisition, videodisc, and CD-ROM control software. Film for the videodisc has been shot during a number of international games and associated training sessions.

Knowledge Processes in Society, Manufacturing, and Human–Computer Interaction

A number of fundamental studies have been undertaken of the way in which knowledge develops in society through a variety of social and psychological processes [5, 7]. The results of these have direct implications for the design of expert systems to support the formation of expertise and to avoid problems of legal liability. The knowledge processes involved in manufacturing have been studied and modeled as part of the development of expert systems for computer-integrated manufacturing. The knowledge structures underlying human–computer interaction have also been studied in depth as part of the development of knowledge-based user interface management systems. We see major advances in a number of disciplines being possible through the development of overt knowledge structures available on shared knowledge bases to scientific and technological groups and communities. Our work in this area is now focusing on computer support of collaborative research, knowledge-based heterogeneous system integration, and the use of object-oriented knowledge bases to support creative thinking processes.

Knowledge Acquisition Workshops

The KSI has been heavily involved in organizing a series of workshops on knowledge acquisition for knowledge-based systems (KAW) attended by research workers worldwide. Three KAWs have been held in Banff, sponsored by the American Association for Artificial Intelligence and co-chaired by John Boose of Boeing Advanced Technology Center and Brian Gaines of the KSI. For the first KAW in November 1986, the intention was to hold a discussion-intensive meeting of some 35 highly involved researchers. In practice, over 120 papers were submitted and some 500 applications to attend were received from some 30 countries. This level of interest lead to a further AAAI KAW in Banff in October 1987 and another in November 1988, to an IEEE-sponsored KAW in England in September 1987, to an

ACM-sponsored KAW in Germany in June 1988, and to a shorter KAW emphasizing integration issues at AAAI in St. Paul in August 1988. Further Knowledge Acquisition Workshops are already scheduled for Paris, Amsterdam, Tokyo, and Banff in 1989 and 1990. These workshops have had a great impact on the literature on knowledge acquisition, resulting in some 60 journal papers and three books to date, and the discussion resulting from them has stimulated knowledge acquisition research worldwide.

Knowledge Engineering Training Program

An intensive one-year training program has been developed for knowledge engineers. Funded by a grant from Employment and Immigration Canada, the program was given to 11 professionals at graduate level from April 1987 through March 1988. The students, with degrees in a range of topics including computer science and psychology, were also employed as research assistants for a number of projects. Courses have been given in knowledge engineering, instructional design, computer hardware and telecommunications, software development and advanced software systems, videodisc development, and human performance.

The training program has aroused much industrial interest and it is expected that the material will be widely disseminated and used for internal training programs and short courses. Arrangements are being made to repeat the program this year.

Research in Machine Learning and Programming by Example

A strong research interest in machine learning has emerged like a phoenix rising from the funeral pyre of adaptive control systems in engineering (some of us have murky pasts). The phenomenon of adaptivity underlies much of our work. So does a central fascination with human interfaces and the communication of information between man and machine. Thus we stress teaching as much as learning; the provision of a congenial environment for casual users to transfer their knowledge and skill.

Robot Programming by Example

Studies are being undertaken to improve communication of human requirements to robot systems. The aim is to enable computer-controlled systems to be taught everyday tasks by people skilled mainly at the tasks. The people should not need, and will not have, the kind of expertise of today's programmers. Existing computer controlled systems — for example, robots and office systems — lack instructability and adaptability. People find it difficult to translate their task skill into explicit procedures for robot, office, and other systems. The research aims to enable non-programmers to specify procedures, without having to do this explicit translation.

A high-precision robot arm has been acquired for studies of robot instructability. Initially the robot will be used for research into techniques for non-programmers to specify tasks to robots, as applied to hospital drug dispensing. It will be tested as part of the Calgary Foothills Hospital automatic robot-based drug dispensing system. The system will select unit doses of drugs and put them into individual patient trays. In a small operation such as this, the pharmacist himself will have to be capable of setting up the dispensing tasks and modifying the robot's performance.

A number of existing programming-by-example systems have been studied and reconstructed in the laboratory. An MIT scheme for creating robot programs from example traces has been reimplemented. A number of shortcomings which were not apparent in the original research for learning object transforms from a user's physical positioning of the arm were addressed. A partially complete program leaves undefined those transforms not available from standard CAD data for a task, and the user can conveniently show the required movements at run time. In addition, the use of quaternions is being investigated as a compact and coordinate independent representation for robot limb rotations, suitable for inducing functional relationships in tasks comprising straight line motions.

Programming by Example in the Human Interface

The user interface provides a multitude of opportunities for programming by example which are being investigated. Rather than concentrating solely on technical methods of inducing programs from examples — an endeavour which became rather sterile in the mid-1980's after an initial flush of enthusiasm — this research emphasizes the central role that interaction can play in clarifying users' intentions. We have studied the acquisition of procedural knowledge, or "know-how," from end users in the domain of interactive graphics. Experts express procedures constructively, using any of the tools available in the interactive drawing environment, and well-structured procedures are inferred using a variety of weak generalization heuristics. A second application domain being investigated is that of casual user interaction with office computer systems using the metaphor of an automated office clerk, while a third is in the realm of text editing. Another example of this kind of work is the "reactive keyboard," invented some years ago at the University of Calgary to help disabled people interact textually with computers.

A central theme underlying this work is the importance of the teacher's view of the learner and his capabilities. Teacher and learner share a common world and have the same capacity for action, an "equal-opportunity" situation. By eagerly carrying out predicted actions (which the teacher can undo if incorrect) the learner may encourage the teacher to satisfy appropriate felicity conditions and make explicit the conditions underlying his actions. In case of ambiguity, dialogue boxes serve to allow teachers to clarify their intentions.

Machine Learning

Many more technical topics in machine learning are being pursued. Our methodology is to begin with reconstruction of influential learning systems which were developed elsewhere. We have reconstructed well-known version space systems, example-generating systems, statistical concept learners, explanation-based learning algorithms, planners, the exploratory learning system AM, and numerous schemes for function induction. In so doing we have learned important lessons and built a springboard for developing our own ideas [13, 25].

Methods for creating structural descriptions from examples involve searching a huge space of possible structural descriptions to check consistency with incoming examples. Efficient ways of representing such spaces in terms of their boundaries are being studied and extended to deal with infinite sets of descriptions (which often arise in practice). Ways of automatically constructing crucial examples which discriminate effectively between competing hypotheses are also under investigation. Shortcomings have been discovered in existing systems for example generation (in particular, the examples generated may be insufficiently general); these are being rectified in a new program which has been built.

Explanation-based learning (EBL) is a semantically based, knowledge-intensive paradigm for machine learning which contrasts sharply with syntactic or "similarity-based" approaches mentioned above. The EBL technique has been redeveloped from the perspective of problem-solving, and studies have been made on how to embed a learning capability into standard inference engines. The work is being extended to the domain of single- and multi-agent planning. We reject the current trend to isolate learning from other activity and study it separately, preferring instead to integrate it into the very heart of the problem solving [12].

Function induction attempts to synthesize a functional relationship from input–output examples. The data may be structured, with numeric and nonnumeric parts. Primitive functions are provided to the system, as are known constants which may be used in functions. We are investigating a number of different algorithms for function induction and comparing their performance both theoretically and practically.

Research in System Theory and Organization

Our fascination with adaptation, mentioned above, also manifests itself lower on the phylogenetic scale of "artificial intelligence" in research on text modeling, prediction, and compression. The representation and management of uncertainty is another common thread which runs through many projects. Also, a strong interest in autonomous and parallel systems colours much of our work.

Text Compression

Complementing the structural work on machine learning is a long-standing and intensive research effort in adaptive text analysis and compression, which grew out of earlier work on adaptive modeling. A text compression method developed at the University of Calgary in 1982–83 still outperforms all known techniques, most of them by a considerable margin. It typically achieves 2.2 bits/character on mixed-case English, and in a recent Australian competition for a compression program both joint winners used variants of our method. A wide variety of compression techniques have been studied and previously unsuspected relationships between them have been found. The two dominant techniques for adaptive coding are statistical and macro methods, and these have been compared and related. Moreover, a very extensive collection of previously described compression methods has been implemented, tested, and assessed [24, 1].

Good compression requires good prediction of text, and studies are being made of the application of the predictive methods to various areas within both the human–computer interface and machine learning; this relates to the programming-by-example research mentioned above [23].

Modeling Uncertain Data

As our systems move more and more into the real world, uncertainty management becomes increasingly important. Rapid prototyping of expert systems involved inferring relationships from uncertain or "fuzzy" data, and the KITTEN system mentioned earlier contains a novel mechanism for fuzzy inference [11]. Within machine learning, efforts are being made to address the problems of creating learning systems that can recover from errors, for it is unlikely that they will find much practical application until they are robust enough to cope with the inevitable vagaries of real-life situations.

Theoretical work on uncertainty management includes a new, unifying framework that encompasses both probability and fuzzy logic, as well as two-valued classical logic, as special cases of a general "standard uncertainty logic" based on valuation of lattice. A separate project has considered the problem of acquiring uncertain rules from examples within a three-valued logic under the ambit of logic programming [2]. A firm theoretical foundation for three-valued logic programming has been established, including a Horn clause subset, procedural interpretation, extension of the standard Herbrand universe and least-fixed-point semantics, and a proof of the correctness of SLD-resolution. Finally, the use of stochastic processes in neural nets has been investigated as a basis for representing uncertain inference.

Autonomous Systems

It is apparent that while many future automated systems will have to operate autonomously, little progress is being made towards this goal. We are looking at the phenomenon of autonomy from three directions [21]. The first regards autonomy as self-government or self-motivation. This is represented by some AI research on representing and using goals, together with physiological, psychological, and philosophical accounts of motivation and goal-seeking behavior. The second concerns the biological independence of organisms which have the ability to maintain their own organization in a capricious environment. The third aims to construct instructible systems that can learn and carry out instructions autonomously but are not self-motivated.

Work in self-motivation is related to the machine learning research discussed previously, in particular the reconstruction of AM. The most concrete examples of artificially autonomous systems of the second genre are various kinds of "deviant" software, worms and viruses, and we have performed experiments and undertaken extensive studies of these phenomena [22]. The third kind includes programming-by-example systems, where the user interface gives learning a human-like autonomy.

Connectionist Architectures

The renewed interest in neural-like processing systems has stimulated a re-evaluation of older paradigms for system organization in terms of massively parallel architectures. For example, we have shown how neural-like systems can be related closely to other well-known formulations of computing such as production systems, finite-state automata, Markov decision processes, and Turing machines. We have studied how novel information representation schemes such as stochastic representations are eminently suited to massive parallelism; benefits of this approach are both theoretical and practical. We have made a detailed theoretical and empirical study of the Boltzmann machine from stochastic, Markovian, and Bayesian viewpoints. We have designed a new VLSI architecture which serves as a building-block for connectionist systems and shown how this can be used to tackle such apparently sequential problems as the n-Queens puzzle and chess playing.

Research in Logic Programming

A major part of the research centers around logic programming, and theoretical and practical studies of programming in logic have been undertaken. (However, we have catholic tastes in language and our implementations are based on

Prolog, LISP — mostly Scheme, and Smalltalk in approximately equal proportions).

Logical Arithmetic

In the past, implementations of real arithmetic within logic programming have been nonlogical. This is a fundamental deficiency which detracts seriously from the logical paradigm. The meaning of programs depends on the order in which statements are executed and on the details of the computer's handling of finite-precision arithmetic. We have discovered how a purely logical interpretation of arithmetic can be created by the use of interval arithmetic as the underlying representation [3]. This can be translated to an implementation within Prolog and invites a redefinition of the language's arithmetic capabilities to use intervals to represent numbers. As well as having a sound logical basis, the resulting system allows a very concise and powerful programming style and is potentially very efficient.

Parallel Prolog using Time Warp

Current research on parallelism in logic programming focuses on AND-parallelism, in which the goals of a conjunction are executed concurrently (the alternative is OR-parallelism which explores different ways of achieving a goal in parallel). The most difficult problem in implementing a distributed AND-parallel system is resolving conflicting variable bindings. Our scheme uses the ideas of "time warp," a mechanism first proposed to implement the virtual time paradigm in distributed simulation, to solve this problem. Simply put, bindings are given timestamps, and conflicts are resolved by comparing these timestamps.

AND-parallel implementations not using the virtual time approach either restrict the parallelism available, or adopt a "committed-choice" strategy and thereby sacrifice the full nondeterminism of standard Prolog. In contrast, our implementation preserves the semantics of sequential Prolog while allowing maximum parallelism. Two considerations motivate this choice: first, for many problems nondeterminism must be retained; and second, maintaining full compatibility with sequential Prolog allows application programs to be ported directly. Variants of the scheme which support nonterminating programs and provide smart backtracking in a sequential system are also envisioned. Another avenue of investigation is adapting Prolog for use in distributed simulation. This adaptation is also based on time warp, but requires parallelism and interprocess communication to be expressed explicitly.

Prolog Programming Environments

We also have a practical interest in constructing programming environments for Prolog. We have built interactive interfaces to the EMACS text editor, a system which displays debugging information graphically [4], a LINT-style utility that checks the internal consistency of Prolog programs, a library facility that allows easy construction of multimodule programs, documentation and typesetting aids in the style of Knuth's "literate programming," and are investigating the practical application of type inference schemes. We have also constructed a number of Prolog interpreters.

Facilities

The Computer Science Department in general, and the Knowledge Science Institute in particular, is well equipped for research. The main departmental workhorses are several Sun-4 systems running Unix (one dual-processor VAX-780 remains as a relic from the olden days). There are a number of high-performance workstations (Symbolics and Xerox LISP machines for AI, IRISs for graphics) as well as many general-purpose Sun workstations. A network of Apollo workstations resides in the KSI, together with a robot arm and the usual assortment of Macintoshes and IBM PC's. A BBN Butterfly is used for work on parallel algorithms. We have access to sophisticated helmet-mounted eye-tracking equipment in the Faculty of Physical Education. The university possesses super-computing facilities (Cyber 205) and supports a large Honeywell Multics operation.

Software resources include Prolog, Scheme, and Smalltalk implementations on most machines (including, for many researchers, Macintoshes at home). The Nexpert expert system shell is used for expert system prototyping and development. Important components of our support environment include Gnu EMACS, TeX, and electronic mail. Virtually all computers are fully networked with high-speed (10 Mbits/s) links. In the summer of 1988 most large departmental research projects, including the AI research, are moving to newly available space, which is being specially modified for our needs, and will be extensively wired with network and other communication facilities.

Researchers

Computer Science faculty engaged in this research include John Cleary, Brian Gaines (the founder and director of the KSI), Paul Kwok, Bruce MacDonald, Mildred Shaw, and Ian Witten. Faculty from other departments who are associated with the KSI are Bruce Clark, Institute for Computer-Assisted Learning; Larry Katz, Faculty of Physical Education; Doug Norrie, Department of Mechanical Engineering; and Joan Vickers, Faculty of Physical Education. Notable visitors who have spent extended periods in the KSI recently include John Andreae, University of Canterbury, New Zealand; Tim Bell, University of Canterbury, New Zealand; Harold Thimbleby, University of York, UK; and Brian Woodward, Knowledge Science Institute.

References

1. Bell, T.C., Cleary, J.G. and Witten, I.H. *Text Compression*, Prentice Hall, Englewood Cliffs, NJ, in press.

2. Cleary, J.G. "Acquisition of uncertain rules in a probabilistic logic" *International Journal of Man–Machine Studies,* 27(2), 1987, pp. 145–154.

3. Cleary, J.G. "Logical Arithmetic" *Future Computing Systems,* 2(2), 1987, pp. 125–149.

4. Dewar, A. and Cleary, J.G. "Graphical display of complex information within a Prolog debugger" *International Journal of Man–Machine Studies,* 25(5), 1986, pp. 503–521.

5. Gaines, B.R. "Foundations of knowledge engineering" *Research and Development in Expert Systems III,* Bramer, M.A. (editor), University Press, Cambridge, 1986, pp. 13–24.

6. Gaines, B.R. "An overview of knowledge acquisition and transfer" *International Journal of Man–Machine Studies,* 26(4), 1987, pp. 453–472.

7. Gaines, B.R. "A conceptual framework for person–computer interaction in distributed systems" *IEEE Transactions on Systems, Man & Cybernetics,* in press.

8. Gaines, B.R. "Knowledge acquisition systems for rapid prototyping of expert systems" *INFOR,* in press.

9. Gaines, B.R. and Shaw, M.L.G. "From timesharing to the sixth generation: the development of human–computer interaction part I" *International Journal of Man–Machine Studies,* 1(1), 1986, pp. 127.

10. Gaines, B.R. and Shaw, M.L.G. "Foundations of dialog engineering: the development of human–computer interaction part II" *International Journal of Man–Machine Studies,* 1(2), 1986, pp. 101–123.

11. Gaines, B.R. and Shaw, M.L.G. "Induction of inference rules for expert systems" *Fuzzy Sets and Systems,* 18(3), 1986, pp. 315–328.

12. Krawchuk, B.J. and Witten, I.H. "Explanation-based learning: its role in problem-solving" *Journal of Experimental and Theoretical Artificial Intelligence,* in press.

13. MacDonald, B.A. and Witten, I.H. "A framework for knowledge acquisition through techniques of concept learning" *IEEE Transactions on Systems, Man, and Cybernetics* (special issue on knowledge engineering), in press.

14. Shaw, M.L.G. "Interactive elicitation and exchange of knowledge in group problem solving" *International Journal of Personal Construct Psychology,* 1987.

15. Shaw, M.L.G. "An Interactive knowledge-based system for group problem solving" *IEEE Transactions on Systems, Man, and Cybernetics,* 1987.

16. Shaw, M.L.G. and Gaines, B.R. "Interactive elicitation of knowledge from experts" *Future Computing Systems,* 1(2), 1986, pp. 151–190.

17. Shaw, M.L.G. and Gaines, B.R. "Advances in interactive knowledge engineering" *Research and Development in Expert Systems III,* Bramer, M.A. (editor), University Press, Cambridge, 1986, pp. 111–122.

18. Shaw, M.L.G. and Gaines, B.R. "KITTEN: Knowledge initiation and transfer tools for experts and novices" *International Journal of Man–Machine Studies,* 27(3), 1987, pp. 251–280.

19. Vickers, J.N. *Instructional Design in Sport and Physical Activity: A Knowledge*

Structures Approach, Human Kinetics, Champaign, Illinois, 1987.

20. Vickers, J.N. and Gaines, B.R. "A Comparison of Books and Hypermedia for Knowledge-based Sports Coaching" *Microcomputers for Information Management*, 5(1), 1988, pp. 29–44.

21. Witten, I.H. "In search of 'autonomy'." *Future Computing Systems*, 2(1), 1987, pp. 35–53.

22. Witten, I.H. "Computer (in)security-infiltrating open systems" *Abacus*, 4(4), 1987, pp. 625.

23. Witten, I.H. and Cleary, J.G. "Foretelling the future by adaptive modeling" *Abacus*, 3(3), 1986, pp. 16–36.

24. Witten, I.H., Neal, R., and Cleary, J.G. "Arithmetic coding for data compression" *Communications of the Association for Computing Machinery*, 30(6), 1987, pp. 520–540.

25. Witten, I.H. and MacDonald, B.A. "Using concept learning for knowledge acquisition" *International Journal of Man–Machine Studies*, in press.

Recent AI Activities at the University of Guelph[1]

By Mary McLeish, Amelia Fong Lochovsky, and David K.Y. Chiu[2]

This report outlines the activities of three groups working in AI related areas within the Department of Computing and Information Science at the University of Guelph. David Chiu discusses a project in pattern analysis and machine intelligence. Amelia Fong Lochovsky discusses several projects in image processing. Mary McLeish reports on her work on uncertainty management and medical diagnostic systems as well as several other projects. The department recently began an M.Sc. program (fall '86), with a particular mandate for applied and interdisciplinary topics.

Image Processing and Computer Vision Algorithms (Lochovsky)

Currently, there is much research attention in parallel architectures for image processing applications. In particular, pipeline architectures have become "extremely relevant for machine vision applications" such as in manufacturing inspection systems [11]. Traditional algorithms often do not take advantage of the inherent parallelism of the hardware. Recently, many important new algorithms have been developed for pipeline architectures [12]. Papers have been published describing various fast algorithms on pipeline architectures and a method has been proposed for analyzing both time and space complexities for algorithms on such architectures [1, 3, 4, 5, 6]. Other research in this area includes work on component labeling algorithms [2].

Image Processing and Pattern Recognition Software Environment (Lochovsky)

Pipeline image processing systems are increasingly being used in industrial vision systems, in particular in manufacturing visual inspection. There has been little work done on the study of the appropriate software environment for rapid program development for these machines. There are general purpose pipeline image processing systems available commercially and they provide a readily available environment for developing new inspection applications before the design of special hardware is undertaken. However, most of these image processing systems do not have the software programming support necessary for a research environment. They provide precanned subroutine libraries which are inadequate when new algorithms have to be developed. Hence, researchers often have to develop the algorithms using low-level machine specific code, which is hard to debug and also difficult to maintain.

Work at the University of Guelph in this area involves the design of a high-level language environment providing software features which reflect the image processing applications as well as the hardware capabilities. Hence the user can program and interact with the powerful hardware in a much easier fashion. Abstraction of various hardware resources such as image storage capabilities and bus connections have been studied. A preprocessor was designed to support the high-level programming environment and the abstract data structures. The preprocessor also analyzes the image processing tasks expressed as expressions and generates the hardware setup and connections necessary for multipass operations through the pipeline processor. A prototype preprocessor has been implemented [7, 8, 9].

Applications (Lochovsky)

Image processing and pattern recognition techniques have been applied to electron microscopic images. This is joint research work with Dr. G. Harauz of the Department of Molecular Biology and Genetics at the University of Guelph. Techniques and software (on the IRIS graphic workstation) have been developed for automatic detection of ribosomes. Ongoing research involves the detection of ribosomes with special characteristics and the recognition of properties for classification purposes [10].

Pattern Analysis and Machine Intelligence (Chiu)

This project concentrates on developing techniques for extracting useful information from empirical data. The study is theoretical and methodological with real life applications and uses methods from mathematics, statistics, information theory, pattern recognition, and artificial intelligence. The work touches on issues in inductive learning, inference, and knowledge acquisition. We have been examining two real-life problems in the research.

---

[1] Reprinted with permission from *Canadian Artificial Intelligence*, July: 18–22; 1989.

[2] Dr. Mary McLeish is an Associate Professor in the Departments of Computing and Information Science and Statistics. Dr. Amelia Fong Lochovsky is currently an Associate Professor in the Department of Computing and Information Science at the University of Guelph. Dr. David Chiu is an Assistant Professor in the Department of Computing and Information Sciences at the University of Guelph.

The first problem involves analysis/synthesis of patterns from molecular sequences. This includes identifying signature subsequences and phylogenetic structure for classification and inferencing. rRNA and tRNA sequences collected from a computer databank have been used. Currently, we are developing algorithms based on the random graph model which incorporates structural and probabilistic variation in the analysis. Other approaches using maximum entropy and event covering have also been evaluated.

The second problem is related to the first but incorporates time as a factor. Historical earthquake data in China have been used to study the phenomenon of earthquakes for long-term prediction. Because of the complexity of the phenomenon, a general theory of earthquakes which yields a practical model of explanation is still missing. We have adopted the pattern recognition approach to evaluate the phenomenon and model it as a system which evolves in time. We have performed several feasibility studies and have gathered some preliminary results.

Theoretical Issues in Uncertainty Management (McLeish)

This work investigates a number of ideas related to probabilistic logic as formulated by Nils Nilsson. The papers [16, 18] study nonmonotonic reasoning in this context. This model allows for inconsistent vectors in Nilsson's model — thus capturing the flavor of default reasoning where it might be considered true that "all birds fly" and that "Tweety is a bird," but the logical or consistent conclusion that "Tweety flies" does not necessarily hold (e.g., Tweety is a penguin). Several ways of incorporating these vectors are considered and the implications for the max entropy and projection solutions are studied. Some comments are also made concerning computational problems.

A later paper [17] develops a relationship between Dempster-Shafer theory and probabilistic logic. In this relationship, the subsets of the outcome set on which belief is assessed in Dempster-Shafer theory are related to sentences in the following manner. The outcome space becomes the set of possible worlds and those worlds in which the sentence is true form the subset of the outcome space representing that sentence. Using the matrix notation of Nilsson and the associated probability vector, the vector V representing the probabilities of the sentences being true become exactly the belief functions Bel, of Dempster-Shafer theory. Instead of the usual Dempster-Shafer theory rules for belief function combination, the entailment schemes of Nilsson are employed. One important difference is that now the probability vector P is not over the set of possible worlds but is over its power set. It is shown that meaningful results can be obtained from the maximum entropy solutions which are defined everywhere for both monotonic and nonmonotonic models. The results are compared with those which would have been obtained from Nilsson's original model and with ordinary Dempster-Shafer belief function combination rules. The results in the case of the appropriate nonmonotonic reasoning model were found to be more realistic than those found in the previous papers.

Another more recent work concerns the answer to a question first posed more formally in the Stanford technical report by Nilsson on probabilistic logic. Two examples are given using the projection and maximum entropy methods of computing a specific (noninterval) entailment value. In one case the answers from the two methods are the same and in the other they differ. Nilsson makes a conjecture as to what might be, at least partially, the cause for the difference in this example and poses the question of what are the general conditions for equality. The paper [24] examines this issue through a study of generalized matrix inverses and gives the family of solutions to the underconstrained system. The paper provides a definite way of testing for equality of the two approaches through a row reduction of the sentence matrix. It also provides an alternate way to find the two solutions and some further examples.

Related to the general problem of nonmonotonic reasoning is work currently submitted for IJCAI 89 [15]. The work by Pearl *et al.* on a calculus of extreme probabilities is considered in some different contexts. Probabilistic logic is investigated for its ability to deal with very small or large probability values and especially in the manner in which they arise in modeling default reasoning. It is again found to be impossible to obtain defined results in probabilistic logic if only consistent vectors are used. If appropriate beliefs are extreme, results can be obtained which give the desired results for nonmonotonic reasoning. In the traditional "penguin" problem, the probability that Tweety flies is still of order epsilon when Tweety is a penguin but the probability that all birds fly is within epsilon of 1. The "ad hoc" Mycin formulas (later ones) for evidence combination are also investigated for extreme values and some interesting results found using odd's ratios.

Applied Project in Medical Diagnostic Systems (McLeish)

A project was initiated in 1987 by members of the Department of Computing and Information Science and the Ontario College of Veterinary Medicine to explore the possibilities for a medical expert system. The OCV hospital has been running a medical information system for several years which collects on-line medical data such as bacteriology, clinical pathology, parasitology, and radiology often electronically generated by lab equipment. Other patient information is also stored such as age, sex, breed, presenting complaint, treatment procedures, diagnosis, and outcome. The intention of the project was to make use of the data to aid the diagnostic process, rather than developing a system based largely on expert opinions. Such an approach has been suggested by Drs. Patil, Szolovits, and Schwarz when they wrote (*New England Journal of Medicine,* 1987) that recent medical systems have been largely based on the assumption that to have expert capability, they must somehow mimic the behavior of experts. "Earlier work using mathematical formalisms (decision analysis, pattern matching, etc.) were largely discarded... ." They suggest it is time to link the old with the new: "now that much of the AI community has turned to causal, pathophysiologic reasoning, it has become apparent that some of the earlier discarded strategies may have important value in enhancing the performance of new programs." The authors recognize the difficulty of this approach when they state that "an extensive research effort is required before all these techniques can be incorporated into a single program."

We have begun by considering two prototype domains. The first concerns equine colic, a complaint of horses often requiring emergency surgery. The decision of whether or not to operate is a difficult one, based on many parameters. A diagnostic chart is now used as a guideline in which occurrence symptoms in two columns are simply added to obtain some idea of the advisability of surgery. A system based on this was implemented in Q'Nial but found to often give equivocal results and produced poor overall performance.

Data were available on 253 cases involving 50 parameters. Outcome information of the following type was available: whether surgery was performed, whether or not a surgical lesion was actually found, and the final state of the animal (i.e., lived, died or euthanized). The data included a number of variable types (real, discrete, nominal) and missing values. Discriminant analysis and logistic regression both found some of the most significant variables to be ones which were very subjectively measured (pain, distension, etc.) and little information was discovered from the clinical pathology data. A PLS1 learning algorithm (due to L. Rendell) discovered some other important variables in the pathology data, which turned out to be most predictive for the outcome information: lived/died. Difficulties were found with Quinlan's algorithm, due to missing data and mixed data types. However, our data have been sent to Australia to be tried on very recent versions of their methodology. Bayesian inductive inference (P. Cheeseman) produced interesting classes in the data, which were often useful in determining special situations in which surgery was indicated but perhaps the overall prognosis was bad anyway and euthanasia the more relevant solution. The results of these studies are given in [21, 22].

Another approach, which has proven quite successful, has been fully implemented in two versions, one using Q'Nial and one using C (in conjunction with Oracle). It involves a weight of evidence combination rule due to I. Good. The continuous variables have been essentially discretized using a fuzzy approach. The clinicians have provided membership functions and appropriate a-level cutoffs have been obtained from the data which maximizes a likelihood function. The evidence combination scheme provides a probability value indicating a level at which surgery is required.

Like many of the probability schemes used, strong independence assumptions are made. In an attempt to capture some of the variable interactions, methods which considered symptom groups, rather than individual symptoms, have been developed. Several search techniques are used to find a set of independent (pair wise) symptom groups of maximal weights of evidence covering an incoming case. This is a tricky procedure involving a balance between several factors, such as small sample sizes and error measurements. Three different heuristic methods have been compared and one selected as the best candidate so far as judged by performance illustrated on roc curve graphs. See references [23, 26, 13] for details.

The prototype involves about 75 symptoms (with the continuous variables discretized) making the computational complexity of the problem very high. It is implemented on a Sequent parallel machine with eight CPUs, and parallel algorithms have been developed and implemented. The machine was purchased through an NSERC equipment grant award (M. McLeish, principal investigator). Other special grants have been obtained from competitions for Excellence Funds through the University of Guelph and NSERC operating grants. A number of other grants have been and are being applied for including the Networks of Centres of Excellence with John Mylopoulos and PRECARN.

A fuzzy approach has now been implemented based partly on the work of the Adlassnig group in Austria. A combination of C and Oracle provides a nice way to implement a fuzzy relational database. The type of sup-min combination rules needed to combine fuzzy relations can be accomplished within Oracle itself, although there are some situations in which invoking C improves efficiency. Further discussions of the details of redesigning the original hierarchical database into a relational model and the construction of the fuzzy relational model can be found in [20]. A later paper [25] describes a blackboard architecture approach to handling the various components of the system. Such a representation seems particularly suitable for a parallel implementation. Work on the colic domain has been conducted with graduate student Matt Cecile, Dr. P. Pascoe of OVC, and Alex Lopez of CIS.

A new domain is now under investigation by a graduate student, Pauline Yao, in conjunction with Dr. T. Stirtzinger, Clinical Studies. This concerns liver disease in small animals. A Dempster-Shafer approach is being tried, but with the difference that the usual probability mass function is being estimated from a study of the data. This has advantages similar to the weight of evidence methodology in that expert opinions can be used to modify the resulting belief functions. (The earlier methodology is being compared with this technique.) The liver disease domain differs considerably from the other prototype in that there is a very large outcome class and symptom group set. The size of the available data sets is too small to rely entirely on traditional statistical and machine learning methods. A linear C code implementation is in place and a parallel algorithm is being designed for this usually exponential method. Several new graduate students have now joined the project.

Other Projects (McLeish)

An earlier write-up in *Canadian Artificial Intelligence* described the rather unusual application domain investigated by Karl Langton. This concerned the Marine Mammal Laboratory in Hawaii under the direction of L. Hermann investigating dolphin communication. Karl brought a number of AI tools to this lab, especially user interface facilities. A natural language system was developed to handle the dolphin's learned vocabulary and grammar. Some of the results of this work appear in [19, 14].

William Hunt has developed an interesting knowledge representation and learning scheme for metaphorical reasoning. The primary form of representation is a hierarchical thesaurus-based approach which includes three units called gestalts, schemata (modeled as frames), and concepts (see work by G. Lakoff). Heuristic search methods are used in the generation process. Using ART and Sun Common LISP, a program has been developed to generate the metaphorical representations. A testing methodology uses human subjects in both the knowledge acquisition of concepts and the evaluation of the computer-generated metaphorical representations and the results have been analyzed statistically. The methodology is able to learn from the source salient features of the target and is designed for everyday phrases and not just scientific analogies. For example, "books are treasure chests" (target: book, source: treasure chest) or "the wind is a cat."

Mary McLeish is currently on sabbatical leave, visiting Dr. D. Spieglhalter at the Medical Research Council in Cambridge, England. Dr. Spiegelhalter's work on the use of Bayesian networks and directed graphs for expert systems has been reported recently in a discussion paper in JRSS, Series B50, 1988. Actual working systems have been implemented in LISP using object-oriented programming. A project is under way using Bayesian networks to model the effects of time in dealing with adverse drug reactions. Dr. Spiegelhalter will be giving a tutorial at IJCAI'89 on the use of Bayesian methods in expert systems.

References

1. Fong Lochovsky, A., "Algorithms and Architectures for a Class of Non-linear Hybrid Filters using Pipeline Architectures," to appear in *Computer Vision, Graphics, and Image Processing*, 1989.
2. Fong Lochovsky, A., "Algorithms for real-time component labeling of images," in *Image and Vision Computing*, Vol. 6, No. 1, Feb., 88, pp. 21–28.
3. Dinstein, I, and Fong Lochovsky, A., "Computing Local Minima and Maxima of Digital Images in Pipeline Image Processing Systems Equipped with Hardware Comparators," *Proceedings of the IEEE*, Mar. 1988, pp. 286–287.
4. Fong Lochovsky, A., "The Effect of Window Mask Shape on Algorithm Design in Pipeline Architectures," submitted to *VISION INTERFACE '89*, sponsored by the Canadian Image Processing & Pattern Recognition Society, June 1989.
5. Fong Lochovsky, A., "Parallel algorithms and Implementation of a class of detail-preserving filters using pipeline architectures, to appear in *Computer Vision and Shape Recognition*, edited by A. Krzyak, T. Kasvand and C. Suen, published by World Scientific Publishing Co. Pte Ltd.
6. Fong Lochovsky, A., "Analysis and Comparison of parallel algorithms for image enhancement in computer vision applications, in *IEEE Computer Society International Computer Science Conference '88 on Artificial Intelligence* Dec. 1988, pp. 473–480.
7. Fong Lochovsky, A., "Implementation Issues of parallel algorithms using pipeline architecture in visual inspection," *Proceedings of the International Assoc. of Pattern Recognition Workshop on Computer Vision*, Tokyo, Oct. 1988, pp. 307–310.
8. Fong (Lochovsky), A., "Increasing Software productivity of Pipeline Image processing Architectures in Manufacture Inspection Applications," *Proc. IEEE 1988 International Conference on Systems, Man & Cybernetics*, Beijing, China, Aug. 1988, pp. 1138–1142.
9. Fong Lochovsky, Amelia C., "High level language abstraction in image processing machines," *Proc. of Second International Symposium on Image Processing, Computer Generated Images, Technology and Applications*, ACM-SIGGRAPH & CESTA, Nice, France, April 1986, Vol. 1, pp. 417–424.
10. Fong Lochovsky, A., and Harauz, G., "Automatic Selection of Ribosomes in Electron Microscopic Images," to be submitted to *Ultramicroscopy.*
11. Sanz, J.L., Dinstein, I., and Petkovic, D., "Computing Multi-colored Polygonals Masks in Pipeline Architectures and its Applications to Automated Visual Inspection," *CACM*, April 1987, pp. 318–329.
12. Sanz, J.L., and Dinstein, I., "Projection-based Geometrical Feature Extraction for Computer Vision Algorithms: Algorithms in Pipeline Architectures," in *IEEE Transaction on Pattern Analysis and Machine Intelligence*, Jan. 1987, pp. 160–168.
13. Gburzynski, P., and McLeish, M., "An Efficient Heuristic Algorithm N-space Partitioning," submitted to the SIAM *Journal of Computing.*
14. Langton, K., McLeish, M., Herman, L., "Delphi: An Intelligent Interface for a Dolphin Communication Laboratory," submitted to the *Second International Conference on Industrial and Eng. Ap. of AI and Expert Systems.*

15. McLeish, M., "Extreme Probabilities and Non-Monotonic Reasoning in the Context of Probabilistic Logic," submitted to *IJCAT '89.*

16. McLeish, M., "Probabilistic Logic: Some Comments and Possible Use for Nonmonotonic Reasoning," in "Uncertainty in Artificial Intelligence," *Machine and Pattern Recognition,* Vol. 5, 1988, pp. 55–63.

17. McLeish, M., "Further Work on Nononotonicity in the Framework of Probabilistic Logic with an Extension to Dempster-Shafer Theory," *Machine Intelligence and Pattern Recognition.* (Edited by T. Levitt and L. Kanal.)

18. McLeish, M., "Dealing with Uncertainty in Nonmonotonic Reasoning," *Proceedings of the Fourth Annual Conference on Intelligent Systems and Machines,* April 1986, pp. 15.

19. Langton, K., Tarbox, B., Herman, L., McLeish, M., "User Modeling at a Dolphin Language Laboratory," *Proc. of the Fourth Conference on Artificial Intelligence Applications,* San Diego, 1988, pp. 400–404.

20. McLeish, M., Cecile, M., and Lopez-Suarez, A., "Database Issues for a Veterinary Medical Expert System," *Proceedings of 4th International Conference on Scientific and Statistical Database Management, Rome,* June 1988, pp. 33–48.

21. McLeish, M., "Exploring Knowledge Acquisition Tools for a Veterinary Medical Expert System," accepted for the *First International Conference on Industrial and Engineering Applications of AI and Expert Systems,* June 1988, pp. 778–788.

22. McLeish, M., "Comparing Knowledge Acquisition and Classical Statistical Methods in the Development of a Veterinary Medical Expert System," *Interface 88 Proceedings* (Statistics and Computing), April 1988, accepted.

23. Cecile, M., McLeish, M., Pascoe, P., Taylor, W., "Induction and Uncertainty Management Techniques Applied to Veterinary Medical Diagnosis," *Proceedings of the AAAI Uncertainty Management Workshop,* Minnesota, August 1988, pp. 38–49.

24. McLeish, M., "A Note on Probabilistic Logic," *Proceedings of the American Association for Artificial Intelligence,* Minnesota, August 1988, pp. 215–219.

25. Cecile, M., Lopez-Suarez, A., McLeish, M., "Using Large Databases of Statistical Information in the Design of a Veterinary Medical Diagnostic System" *Proceedings of the 1988 AAAI Workshop on Databases in Large AI Systems,*" Minnesota, August 1988, pp. 36–46.

26. McLeish, M. (et al.), "Uncertainty Management and Knowledge Acquisition Techniques from a Medical Information System," accepted for the *2nd International Workshop on AI and Statistics* (W. Gale, Bell Labs), Jan. 1989, submitted for *Annals of Math and AI.*

Artificial Intelligence at Queen's[1]

Edited by Janice Glasgow[2]

AI is an active area of teaching and research at Queen's University. With respect to teaching, Queen's offers a specialized degree in Cognitive Science. This program is administered jointly by the Departments of Psychology and Computing and Information Science. It involves breadth of knowledge about psychology and computing with depth in the areas of AI and cognitive psychology. There is also a variety of graduate AI-related courses offered at the university.

Research in AI at Queen's is concentrated in the areas of robotics, perception, knowledge-based systems, and music. Following is a description of some of the current AI projects at Queen's.

Robotics and Perception

Researchers: Department of Computing and Information Science: Roger Browse, Randy Ellis, Dorothea Blostein; Department of Psychology: Susan Lederman, Terry Caelli; Department of Electrical Engineering: Mohammed Bayoumi.

Most existing robotic systems are only capable of operating within a completely understood layout of the workspace. Providing that each situation is a replication of a prototype environment, the robot may perform its actions effectively. Robots of the future will be expected to operate in imprecise and changing environments, and will therefore require sophisticated perceptual capabilities. Artificial intelligence projects within the Queen's Robotics and Perception Laboratories are directed at issues in computational perception and support for intelligent robots.

Terry Caelli and Roger Browse are investigating the problem of image segmentation using techniques such as controlled relaxation and spatiochromatic feature filters derived from supervised learning conditions and separated feature maps as proposed in feature integration theory. Generalized texture discrimination shares many aspects of the problems of segmentation, and so this topic is also addressed, particularly with respect to human performance in conditions of controlled texture differences. In addition, Dorothea Blostein is investigating the use of texture information in the derivation of surface orientation. Browse and Caelli research the questions of object recognition and part identification utilizing both canonical view methods and object model approaches.

In the experimental psychology division of the Robotics and Perception Laboratories at Queen's, Susan Lederman and her collaborator Roberta Natzky (University of California at Santa Barbara) have demonstrated that human recognition of objects by touch is fast, accurate, and flexible. Their work over the past few years constitutes one of the first attempts known to provide a systematic theoretical framework and comprehensive experimental approach aimed at modeling the intelligent exploration and recognition of objects by human touch (with and without vision). They have used a variety of experimental methodologies successfully in their work, all of them involving the analysis of human hand movements and/or other performance measures during human tactile object exploration and identification. Such techniques have permitted them to address the following kinds of questions: Are there clear associations between the way a human manually explores an object and the type of object knowledge required (shape, hardness, texture, etc.)? Since this has been shown to be the case, what is the nature of these links? Why are particular hand movements used to extract particular kinds of information about three-dimensional objects? What factors affect the order in which intelligent exploration and information processing about objects occurs, and how? What is the nature of the representations of objects in memory, and the processes (data-driven and model-driven) by which these representations are created when touch is used (with and without vision).

Another project in the lab addresses the nature of preattentive and attentive processing of stimuli presented to the skin. To continue and expand the initial work done by Lederman, a custom-designed apparatus is currently being constructed, in conjunction with Tom Moore of the Mechanical Engineering department. This equipment, known as Q'HAND for Queen's Haptic Attention Device permits stimuli (e.g., surfaces varying in texture, material, compliance, form, etc.) to be delivered to any of several fingers on each hand under computer control. This has rarely been attempted in studying tactile information processing to date and as such offers an opportunity for exploring these issues in considerably greater depth.

---

[1] Reprinted with permission from *Canadian Artificial Intelligence*, July: 22–25; 1989.

[2] Janice Glasgow is an Associate Professor in the Department of Computing and Information Science at Queen's. Her research interests include programming language semantics and AI.

There are exciting and potentially significant spinoffs of the work on human object perception and recognition by Lederman and Klatzky (described above) in industry. They are attempting to apply knowledge of the structure and function of biological perceptual systems (touch, with and without vision) to the design of flexible perceptual systems for sensate dextrous robot hands, capable of intelligent exploration, perception/recognition, grasping, and manipulation of objects. The first successful robotic application of Lederman and Klatzky's work with humans has been implemented by Sharon Stansfield under the supervision of their collaborator Ruzena Bajcsy, who heads the perceptual robotics GRASP lab at the University of Pennsylvania. Stansfield developed a tactile perceptual system capable of recognizing a considerable number of common objects, using a stereo vision system to locate the object and a Puma 560 arm equipped with a Lord LTS-200 array force sensor on the wrist to actively explore the objects.

An important aspect of this novel collaborative team is the strong emphasis on active sensing and perceptual exploration by sensate robotic systems. This can be used effectively in a variety of situations. For example, the current NASA practice is to replace an entire instrumentation panel of the space shuttle when a single component becomes defective. This is not a cost-efficient procedure, however. An alternate solution is to have a sensate robot actively explore the panel to locate and subsequently replace the faulty part. Carrying into space only small components, as opposed to relatively massive sections such as the entire panel, would conserve valuable space aboard the shuttle. Data-driven robotic exploration can be used to determine the material composition of soil on unfamiliar planets.

Randy Ellis's work in robotics principally involves how to design and use tactile sensors to enhance the perceptual and manipulation capabilities of advanced robotic systems. His current work concentrates on three problems: the propagation and management of uncertainty in perception; the theoretical specification of tactile sensors, and issues in force and tactile feedback for teleoperation systems.

The uncertainty project is driven by the observation that sensor readings, regardless of their source, inherently contain errors. In a model-based recognition system, these errors propagate as uncertainties through the signal processing, initial matching, data interpretation, and possibly into a stage of planning a new sequence of actions to aid in the recognition of the object. Using a computational geometry approach, he has found tight analytic bounds on the uncertainty in a tactile recognition paradigm and has used these results to produce a robotic system which reasons about how to plan the "best next path" for subsequent contacts with an object in a robot workspace.

A frustrating problem in tactile perception is that sensor design and sensor use have been proceeding along divergent paths. Ellis is investigating the possibility of specifying parametric relations between robot limb dynamics, tactile sensor characteristics, and task specifications, with the goal of being able to determine what sensor properties are needed to accomplish a given task on a given robot. Ideally, such specifications could then be used to guide the design of more advanced tactile sensors than are currently available.

One ambitious project of the laboratory is teleoperation. It is a collaborative effort with Browse, Bayoumi, Ellis, and Lederman. The goal of the project is to provide sophisticated, multisensory feedback to a human operator who is guiding a remote robot through its motions. Ellis is investigating the design and control of a force-feedback position controller, which will permit the operator to "feel" the forces acting on the end effector of the robot as the robot imitates the operator's motions. He and Lederman are also investigating issues in the human's interaction with the system, to find convenient and powerful presentations of multisensory information and to characterize the operator's performance in common manipulation tasks.

Knowledge-Based Systems

Researchers: Computing and Information Science: Janice Glasgow, Mike Jenkins; Mechanical Engineering: Brian Surgenor.

The Nial Project at Queen's University is involved in the design, implementation, and application of the Nested Interactive Array Language, Nial. One of the principal application areas of Nial is artificial intelligence, where the powerful operations for dynamic manipulation of both symbolic and numeric data permit straightforward implementation of AI techniques.

Research in the area of knowledge-based systems falls into two categories: tools and applications.

AI Tools

The Nial AI Toolkit is one product of the ongoing research of the Nial Project. The philosophy of this work is to provide building blocks from which tailored knowledge-based systems can be constructed. The goal is to provide direct implementations of basic techniques that can be understood well enough to be modified by the programmer constructing a particular application.

One tool that has been developed is a logic programming environment. The goal of this research was to integrate the functional style of Nial with the declarative capabilities of logic. A variety of theorem provers have been implemented, including ones with semantic unification and heuristic search strategies. These provers vary in both flexibility and efficiency.

Semantic unification replaces the need for built-in predicates in the logic programming system. Since any Nial expression can be a term in the logic, both arrays and function applications on arrays can be represented in the language. The underlying array theory semantics of Nial provides a basis for treating arrays as logical terms in the language. The approach being used is based on an axiomatization of array theory in which arrays are constructed using two fundamental functions of the theory: reshape and hitch.

A major scientific result of research in this project was the development of an environment in which both databases and logic programs can be represented and manipulated. One reason for this success is based on the mathematical formalism of array theory, on which Nial is based. A roster database is an array structure that can be manipulated using functions of array theory. Similarly, logic clauses can be represented as embedded arrays and mathematical functions (unify and resolve) can be expressed in array theory to define control strategies for the logic. Thus a single semantic entity can be used for both forms of representation. The result of this research is the ability to explore the advantages and disadvantages of separating factual information, such as that stored in a conventional database, from knowledge or general rules that relate these facts.

Other tools developed are a frame language and forward and backward rule interpreters.

Applications

The programming language Nial has been used at several universities and industries to develop prototype knowledge-based systems. These include systems for aircraft maintenance (Brunel University, England), well logging (Alberta Research Council), fuzzy sewage and system control (IBM and Technical University of Denmark), automatic software testing (IBM Toronto), and insurance policy screening (two insurance companies). As well, Queen's University is playing an active part in the Computer Aided Process Engineering Lab, a subgroup of the Manufacturing Research Corporation Centre of Excellence for Ontario. In this project we are collaborating with chemical and mechanical engineers from the Universities of Waterloo, Toronto, and Queen's in developing prototype knowledge-based systems for process engineering.

The Nial Project at Queen's is currently developing two prototype expert systems under contract from industry:

- *Automated Data Management System for Mobile Servicing System*: This system is being developed in collaboration with Spectrum Engineering of Peterborough and is funded by the Strategic Technologies for Enhanced Automation and Robotics (STEAR) program. Major tasks of this system include fault diagnosis, sensor data handling, sensor trend analysis, and reliability maintenance and reporting. Because of the need to manage both data and knowledge in this system, Nial has proven to be a useful implementation language.
- *Depreciation Knowledge-based System*: The capital recovery group in Bell Canada have described their requirement for a computer system to assist them in doing life analysis of the company's property. This project involves Bell and the Nial Project at Queen's in a research contract to explore the development of such a system using knowledge-based systems

technology and the programming language Nial. This system will interact with an existing process developed by Bell Canada that automates much of the routine computations required to do life analysis. This new system is expected to make suggestions on what curves and lives best represent the expected mortality based on inputs by the user of past history and factors that will affect future lifetime trends.

- *Control and Operation of a Thermal System*: It has been suggested that AI can provide the means to significantly improve the performance of process operations by providing the operator with timely expert advice on optimal control and procedures for a given mode of operation. Brian Surgenor is currently investigating the performance of a prototype knowledge-based system applied to a physical analog of a superheater steam temperature system.

Artificial Intelligence and Music:

Researchers: Department of Computing and Information Science: Dorothea Blostein, Eli Blevis, Peter Farrett

The work of Peter Farrett involves symbolic logic as applied to knowledge-base modeling. In particularly, he is considering monotonic and nonmonotonic models. The domain of discourse is the "composer's assistant," a tool for music composition which models the behavior of the end user. This is an interdisciplinary field and merges advanced mathematics, musical knowledge, and computer science. It is hoped that this research will extend the capability of knowledge-base system technology in current computer music applications.

Eli Blevis is involved in the construction of a programming language that allows composers to express what they do in a manner that is meaningful to them. This path has led him to consider the construction of the language from a first semantic principle: what is a music logic and how does it differ from logics that are intended to model notions of truth or reality? The importance of this work for computer science is hopefully a demonstration of the influence of modern thought about logics to application-specific programming language design and semantics.

Successful computer vision depends on combining low-level primitive-extraction processes with high-level knowledge-intensive interpretation processes. In order to study the modeling of world knowledge, and the use of these models in image interpretation, it is helpful to focus on restricted image domains that have well-defined semantics.

Dorothea Blostein is investigating methods for extracting the semantics of two-dimensional data representations such as graphs, charts, sheet music, or wiring diagrams. Such diagrams are governed by strict rules and conventions. Since these conventions are well defined in an informal way, research can focus on how to formalize these notational conventions in a representation suitable for image interpretation. For concreteness, she is focusing on a particular type of diagram (sheet music), but with the goal of developing general representation and interpretation techniques which can be adapted to a variety of notations. This problem has important practical applications. For example, a utility company which wishes to computerize its record-keeping may have huge quantities of diagrams (showing locations of pipes and cables, manholes, etc.) which need to be entered into a computer system. This is a tedious and time-consuming task for humans; it is clearly preferable for the computer to be able to interpret images of the data directly.

Current Funding of Described Research

NSERC (Operating, Infrastructure and Equipment); NRC (STEAR Program); Manufacturing Research Corporation of Ontario (MRCO); Information Technology Research Centre (ITRC); U.S. Office of Naval Research (ONR); Bell Canada.

References

1. R. Bajcsy, S. Lederman, R. Klatzky, "Machine systems for exploration and manipulation: A conceptual framework and method of evaluation." Submitted.

2. E.B. Blevis, M.A. Jenkins, "Computing Similarity in Music Presentation" at the *First International Workshop on Artificial Intelligence and Music, Gessellschaft fur Mathematik and Datenuerar-beitung,* St. Augustin, West Germany 1988.

3. E.B. Blevis, M.A. Jenkins, J.I. Masgow, "Motivations, Sources, and Initial Design Ideas for CALM: A Composition Analysis/generation Language for Music." *Proceedings of the*

First Workshop on A.I. and Music at AAAI-88 St. Paul, Minnesota 1988.

4. E. Blevis, "A Basis for Effective Logic Programming in Nial" Master's Thesis Queen's University at Kingston 1986.
5. E.B. Blevis, M.A. Jenkins, E.P. Robinson, "On Seeger's Music Logic" INTERFACE: *Journal of New Music Research.* (To appear, Summer 1989.)
6. D. Blostein, N. Ahuja, "A Multi-Scale Region Detector" Computer Vision, Graphics, and Image Processing Accepted June 24, 1989.
7. D. Blostein, N. Ahuja, "Shape from Texture: Integrating Texture-element Extraction and Surface Estimation," Submitted to *Pattern Analysis and Machine Intelligence.*
8. D. Blostein, N. Ahuja, "Representation and Three-Dimensional Interpretation of Image Texture: An Integrated Approach." *Proceedings First International Conference on Computer Vision,* London, England June 1987.
9. D. Blostein, "Recovering the Orientation of Textured Surfaces in Natural Scenes" Coordinated Science Laboratory, *University of Illinois Technical Report UILU-ENG-87-2219,* 1987.
10. R. Chau, J. Glasgow, M. Jenkins "A Framework for Knowledge Based Systems" in Nial *Proceedings of the 6th Annual IEEE Phoenix Conference on Computers and Communication,* Feb. 1987.
11. R. Chau, "A Model for Knowledge Information Processing Systems" in *Nial Master's Thesis,* Queen's University at Kingston, Jan. 1987.
12. R. Chau, J. Glasgow, M. Jenkins, "Fuzzy Information Management Using the Roster Model." *Proceedings of 21st Hawaii International Conference on System Sciences,* January 1988.
13. R.E. Ellis, "Uncertainty Estimates for Polyhedral Object Recognition" *Proceedings of the IEEE Conference on Robotics* 1989 (in press).
14. R.E. Ellis, "A Tactile Sensing Strategy for Model-based Object Recognition" Computer and Information Science, University of Massachusetts Ph.D. Thesis, published as *COINS Technical Report* 87–96, 1987.
15. R.E. Ellis, E.M. Riseman, A.R. Hanson, "Tactile Recognition by Probing: Identifying a Polygon on a Plan" *Proceedings of the AAAI Conference,* pp. 632–637, 1986.
16. R.E. Ellis, "An Approach to the Integration of Vision and Touch for Robot Control" *Proceedings of the IFAC Symposium on Robot Control,* pp. 343-348, 1985.
17. J. Glasgow, M. Jenkins, C. McCrosky, "User Defined Parallel Control Strategies" *Proceedings of the Second IEEE Symposium on Logic Programming,* 22-28 Boston 1985.
18. J. Glasgow, M. Jenkins, C. McCrosky, H. Meijer, "Expressing Parallel Algorithms in Nial" *Parallel Computing,* Feb. 1988.
19. J. Glasgow, M. Jenkins, "Array Theory, Logic and the Nial Language" *Proceedings of the International Conference on Computer Languages,* Miami, Florida, October 1988.
20. L. Hache, "NFL: A Frame Representation Language in Nial Master's thesis" Queen's University at Kingston 1986 *Nial Frames Thesis.*
21. M. Jenkins, J. Glasgow, "A Logical Basis for Nested Array Data Structures" Queen's University at Kingston *Technical Report # 87-205* (Revised Feb. 1988) December 1987.
22. M. Jenkins, J. Glasgow, C. McCrosky, "Programming Styles in Nial IEEE Software Engineering" 46–55 1986. Also available as *Proceedings of the 19th Hawaii International Conference on Systems Sciences.*
23. M. Jenkins, W. Jenkins, Q'Nial AI Toolkit Nial Systems Ltd., Kingston, Ontario 1987.
24. M. Jenkins, J. Glasgow, E. Blevis, R. Chau, E. Hache, D. Lawson, "The Nial AI Toolkit" *Proceedings of Avignon '88 Eighth International Workshop on Expert Systems and their Applications,* June 1988.
25. R. Klatzky, S. Lederman, C. Reed, "There's more to touch than meets the eye: relative salience of object dimensions for touch with and without vision" *Journal of Experimental Psychology: General* 1987 116(4) 356–369.

26. S. Lederman, R. Klatzky, "Hand movements: A Window into haptic object recognition" *Cognitive Psychology* 1987 19(3) 342–368.

27. S. Lederman, R. Browse, R. Klatzky, "Haptic processing of spatially distributed information" *Perception & Psychophysics*, 1988 44(3) 222-232.

28. S. Lederman, R. Browse, "The physiology and psychophysics of touch" (pp. 71–91). In P. Dario (Ed.), Sensors and sensory systems for advanced robots. NATO ASI Series F: *Computer and Systems Science*, Springer-Verlag: Berlin 1988.

29. S. Lederman, R. Klatzky, "Haptic exploration and object representation." In Goodale, M. (Ed.) *Vision and action: the control of grasping*. New Jersey: Ablex, in press.

University of Saskatchewan

By Gord McCalla, Professor

Introduction

Of 16 full-time faculty members in the Department of Computational Science, five have primary research interests in artificial intelligence. Gord McCalla is involved in a wide variety of AI topics, including representation of knowledge, natural language understanding, dynamic planning, and diagnosis. These interests are currently focused in the area of intelligent tutoring systems. Jim Greer is also interested in intelligent tutoring systems, as well as cognitive modeling and qualitative reasoning. Eric Neufeld does research in foundational areas of knowledge representation with a particular emphasis on probabilistic reasoning. John Cooke has long-standing expertise in cognitive science and a more recent interest in uncertainty management in rule-based expert systems. Herb Yang has been involved for many years in a wide range of image processing research, including shape analysis, segmentation, motion analysis, the application of color, and the physical principles as well as the mathematics of computer vision.

In addition, a number of other Department members have interests in areas closely related to artificial intelligence. Jim Carter carries out research in human factors, especially intelligent, integrated system design methodologies. Paul Tremblay is involved with an IRIS subproject whose aim is to use AI techniques in the design of large-scale information systems. Finally, Marlene Jones of the Alberta Research Council brings her expertise in AI applications in education to the Department in her role as an adjunct professor.

A description of a number of AI projects in the Department follows, examples of the kinds of things currently underway. These projects are funded through a variety of means — Networks of Centres of Excellence funds, NSERC Operating Grants, and University of Saskatchewan Scholarships to many of the graduate students.

Some Projects

Projects in the ARIES Laboratory (McCalla, Greer, Jones)

The Laboratory for Advanced Research in Intelligent Educational Systems (ARIES) was founded in 1987 to provide a focus for the growing number of projects concerned with applying AI techniques to training and education. Artificial intelligence can help traditional computer assisted instruction approaches to be more flexible and more oriented to the student. AI can also help in the design of learning and training environments in which the student can discover ideas for himself or herself.

The major ongoing project in the ARIES Laboratory has been the SCENT (Student Computing Environment) project, which started in the mid-eighties under the direction of Gord McCalla and Rick Bunt, who was subsequently replaced by Jim Greer as a principle scientist. SCENT has as its eventual aim the creation of a supportive and intelligent programming environment for use by novice programmers. Many investigations over the years have looked into various aspects of SCENT, including recognition of student strategies from their code, recognition of misconceived strategies, blackboard control structures for SCENT, and empirical analyses of students' actual program creation techniques. More recently, a new SCENT-3 architecture has been created to guide the next generation of SCENT investigations. Granularity-based recognition, instructional planning, mental models of recursion, student model maintenance, and the design of knowledge engineering tools have all been subjects of current research. Among the graduate students and research assistants working on SCENT over the years have been Janelle Harms, Teng Ng, Paul Pospisil, Bryce Barrie, Xueming Huang, Barb Brecht, Judy Escott, Randy Coulman, Mary Mark, Shawkat Bhuiyan, Luong Nguyen, and Mike Oliphant.

In addition to SCENT, there have been a number of other projects formulated under the auspices of ARIES. One of these, carried out by Dinesh Gadwal, is the design of a knowledge-based chess tutor, resurrecting chess as an interesting domain for AI issues beyond just search. Another ARIES project is being carried out by Dan Baril into the use of qualitative reasoning techniques in tutoring about complex devices. Mary Mark is embarked on a project to investigate how a concern for evaluation can be built into a tutoring system from the beginning, thus enhancing the ability to actually evaluate the system later. Barb Brecht and Marlene Jones have been involved in experiments with the genetic graph as a means of tracking and predicting student learning paths. Marlene Jones is also interested in educational diagnosis and curriculum design. Nascent projects are underway investigating tutorial dialogue (Carl Gutwin), instructional delivery planning (Permanand Mohan), case-based instructional planning (Zhihui Du), tutoring literacy (Maya Kurup), and granularity approaches to explanation-based learning (Srinivas Palthepu).

IRIS Subprojects (Neufeld, Greer, McCalla, Tremblay)

The federal government's Networks of Centres of Excellence project, IRIS, has a number of participants from the Department of Computational Science. Paul Tremblay, Gord McCalla, and Jim Greer are collaborating in a subproject of IRIS, concerned with the design and construction of large-scale information systems. Tremblay is investigating (with ex-colleague Paul Sorenson, now at University of Alberta) the production of meta-facilities for CASE tools; and Greer and McCalla are designing an intelligent help system to aid system users in understanding the often complex information system design tools. Eric Neufeld is involved with another subproject of IRIS concerned with formal knowledge representation. His particular emphasis is on how notions of common sense reasoning can be captured in probabilistic terms.

Image Processing Projects (Yang)

Herb Yang has been involved for a number of years in the area of image processing. A number of projects have been and are being investigated. One such project involves the development of techniques that can be used in the recognition of shapes that are deformed under controlled conditions. Another project looks into human motion understanding, with the goal of developing a usable system for tracking, labeling, and understanding human body motion. Computational color stereopsis is the focus of another investigation, where a new technique for stereopsis incorporating color information seems to be better than existing achromatic approaches. This work has led to an interest in basic issues in color computer vision, and further work is underway to determine factors affecting the use of color in both computer and human vision systems. Practical applications of Herb Yang's vision work are currently being developed in the area of biomedical imaging, in collaboration with the Department of Obstetrics and Gynaecology at the University of Saskatchewan and also with the Cell Tracking Group at McGill University. Over the years many graduate students have been instrumental in the imaging work, including Doug Brockelbank, Maylor Leung, Jian Xu, Warren Long, David Callele, and Sudarsan Tandri.

Uncertainty Management in Expert Systems (Cooke, Neufeld)

Under the supervision of John Cooke and Eric Neufeld, graduate student Kofi Sapong is investigating various approaches to uncertainty in expert systems, including Bayes' theorem, the certainty factor approach, the Dempster-Shafer theory, and fuzzy logic. These approaches are being evaluated by comparison to the decision making of human experts in the field of finance.

Conclusion

As can be seen from the foregoing project descriptions, AI is a vigorous and important part of the Department of Computational Science at the University of Saskatchewan. (The other kind of AI is pursued in the College of Agriculture and will not be discussed here.) Artificial intelligence research should continue to prosper in the coming years, as the more junior researchers gain in status, the more senior researchers consolidate their reputations, and as the many graduate students leaving the program begin to impact their future places of employment.

Simon Fraser University[1]

In 1976, Nick Cercone, a young University of Alberta graduate, was hired to teach a course at Simon Fraser University (SFU) called Computing and the Humanities. An underlying element was that people in the Humanities could benefit from artificial intelligence (AI) research. For instance, it was recognized that literary experts and linguists often need to manually attach semantic or syntactic tags to words in their texts, whereas the possibility existed that this job could be done by a computer with AI.

As Dr. Cercone began attending international conferences on Computing and the Humanities in the late 70s it was soon realized that SFU's emphasis on interdisciplinary studies would be well served by a strong AI presence. Teaching of the basic concepts of AI and computational linguistics commenced. Soon thereafter, Brian Funt, an expert in computational vision systems, was hired, following which the linguist and logic programming expert, Veronica Dahl, joined the SFU AI team. From these beginnings, AI research at SFU has increased to include 11 core professors and several dozen research associates. In addition, there are other professors whose involvement continues to grow. Many SFU faculty have demonstrated their interest in AI, but have not received formal AI training. Examples are Tom Calvert in Kinesiology and Tom Poiker in Geography.

Research at SFU covers five main areas of classic AI: expert systems, computational vision, computational linguistics, knowledge representation, and learning systems. The result is a broad range of topics that are supported here with a balance that could not be achieved if effort were concentrated in any one area.

On the theoretical side, research ranges from the formal representation of knowledge and meaning to logics of cognition and belief. More practical investigations concern the nature of vision, color perception, and motion planning, as well as natural language understanding and machine translation. At the most applied level, work is underway on expert systems at the heart of new powerful computer-aided design systems, while others study the man–machine interface using sophisticated text, graphics, and speech recognition software.

While this brief report details the work of the core researchers, many other professors at SFU are also active in AI. Binay Bhattacharya applies computational geometry to AI problems. The philosopher, Steven Davis, studies cognitive aspects of pragmatics as well as theories of action and belief. Jiawei Han in Computing Science extends his expertise in database systems to methods of knowledge discovery. Meanwhile, Philip Winne in Education is using AI to help teachers improve lesson planning while also working on an intelligent tutoring system.

To many professors, AI research is one of the greatest intellectual challenges of all time. It poses questions like: What is Intelligence? How does knowledge come to have meaning? And what is the nature of Belief? N. Cercone likes to call it a sort of "applied epistemology." If researchers at SFU are not actually discovering the "meaning of life," at least they are working on problems that give life meaning.

AI Specialists and Their Research

Reasoning Engines

The goal of Dr. R. Alleliunas is to show how logic programming, AI, and related ideas can have a significant practical impact, often in unexpected ways. The value of linking research and commercial development is clearly recognized and maintained.

Process control has been identified as an area that can benefit from the difficult marriage of symbolic computation (i.e., AI, Prolog) and realtime applications. One project nearing completion is a custom expert system for realtime process control at Johnson Matthey Electronics. This system turns a friendly high-level process description into low-level device control signals. Johnson Matthey will be using it for the production of gallium arsenide semiconductor crystals.

Dr. Alleliunas's most recent research involves making stochastic models computationally more tractable. Long before the advent of computers, the value of stochastic models was well known to engineers. They are an alternative to typical rule-based expert systems presently in use.

---

[1] Extracted with permission from *Update*, April 1990, Vol. 2, Number 3, "Centre for Systems Science, Simon Fraser University."

Natural Language and Knowledge Representation

In addition to acting as Director of the Centre for Systems Science (CSS), Dr. Cercone's theoretical research includes a representation for natural language understanding using extended semantic networks. The representation used can be formally interpreted (in the Tarskian sense) and does not require special purpose representations at the conceptual level. A computational method for selecting possible antecedents of referring expressions over sentence boundaries is under development. This model uses a λ-categorial language for meaning representation and incorporates valuable features of Fregean-type semantics and situation semantics developed by Barwise and Perry.

Researchers at the CSS are also introducing a new approach to representing and manipulating nonsingular concepts in natural language discourse. This representation systematically distinguishes between those language terms that are used to refer to objects classified within different but related levels of the model and, in addition, captures inter-sentential dependencies in discourse. Also, the representation serves as a basis for defining a formal semantics of discourse paragraphs.

The prototype English-to-SQL translator (natural language interface to relational databases) developed at Simon Fraser has resulted in several major advances including learning of semantic and syntactic English word features through dialogue with the user, computing and utilizing correct navigation (join path) information through usage, and portability across databases.

Future work will involve an expert system database designer and analyzer with a more robust query interface and flexible user interviewer. Design goals include portability, automated knowledge acquisition, user flexibility, and synergistic knowledge management.

Language Translation and Generation

An Associate Professor in Computing Science, Veronica Dahl, is interested in the uses of logic programming for language processing, for deductive database and expert systems, and in the development of logic programming tools to solve problems in these areas.

Current research includes automation of linguistic theories, such as Chomsky's, development of logic grammar tools specifically suited to implement such theories (Discontinuous Grammars in particular), development of unification-based schemes for type hierarchies in deductive databases, and translation of the Static Discontinuity paradigm — one of the Discontinuous Grammar formalisms — into logic programming proper, to be used, for instance, in problems of design.

Current applied research includes the automatic translation of agricultural market reports between French and English.

Formal Knowledge Representation

The research work of Dr. Jim Delgrande involves extensions to first-order logic (FOL) that address forms of defeasible reasoning. For instance, common sense conditionals like "Birds fly" are handled by adding a variable conditional operator to FOL to represent such assertions. Essentially the variable conditional is read as "in the normal course of events, birds fly," so the default reasoning is then paraphrased as "if something is a bird and I have nothing to make me believe otherwise, then that thing flies."

Additional work includes study on the problem of limited and conflicting beliefs. An omniscient agent would know everything, but humans — and computers in particular — simply do not know all the consequences of their beliefs. Embracing contradictory ideas is an essential part of being human: we have all scheduled two meetings for the same time. Work includes the exploration of formal theories of Belief that model an intelligent system that does not know the full set of logical consequents of its beliefs.

Color Constancy and Vision

The research work of Dr. Brian Funt addresses the question of how we see color and, in particular, color constancy — how our color perception remains relatively unchanged under varying lighting conditions. From a practical standpoint, color constancy is very important for us, since the color variations we encounter, say as a cloud passes in front of the sun, are quite dramatic despite the fact that we do not experience them as such. Without color constancy, we would be lost in a world of varying colors. Similarly, a robot depending on color to find its way but without a capability for color constancy would also quickly become lost in the face of such disturbances.

Consideration is therefore given to developing computer algorithms that exhibit color constancy. One approach to the problem is to seek additional sources of information about the light reflected from a surface over and above that provided by the color receptors in the eye. For example, an analysis of the chromatic aberration caused by the eye's lens shows that just such additional information about the spectrum of the reflected light can be extracted from an image. A paper on this work won the 1988 IEEE Marr Prize.

Motion Planning and Machine Vision

The two components of an intelligent robotic system are perception and action. The research efforts of Dr. Kamal Gupta involve both areas. With robot perception, the goal is to develop theory and implement algorithms for obtaining 3-D models of objects in a given scene, both from range and intensity images. This covers multisensor integration and surface reconstruction and analysis.

For robotic action, one seeks to develop efficient geometric algorithms for tasks such as motion planning and grasping. Research efforts by Dr. Gupta in this direction have resulted in efficient motion planners for manipulator arms.

Additional work includes the investigation of efficient algorithms for planning robot motions among moving obstacles in polygonal domains where the research efforts have resulted in a space-time framework for efficient solution of such problems.

Philosophical Issues of Cognition

Knowledge representation and machine learning form the two research interests of Dr. Bob Hadley. Within knowledge representation he is concerned with formal models of belief, the logic of truth-predications, and how belief is related to a cognitively sound theory of semantics. Recently work has focused upon the relation of formal belief models to intensional theories of semantics and has criticized some prominent epistemic logics within AI.

Research in machine learning questions the robustness of neural networks as an adequate model of the acquisition of syntactic and semantic knowledge. Serious problems have been found with connectionism, a widespread approach towards cognitive modeling. One recent paper in this subject area questions the claim that higher level cognitive processes possess an essentially connectionist (or neural network) architecture.

In addition, a research project being started will be concerned with ways of using connectionist networks to implement a particular lower level cognitive process. The project will attempt to simulate the processes by which simple referring words in a child's primitive vocabulary become associated with the objects and actions in the world denoted by these words. The process is far more complex than is often supposed.

Reasoning Systems

Dr. Bill Havens is concerned with the development of constraint reasoning technology for diagnosis, synthesis, and other recognition tasks in addition to being Director of the SFU Expert Systems Laboratory.

In the SFU Expert Systems Laboratory we are developing a new expert system reasoning engine, Echidna, which incorporates constraint logic programming, schema knowledge representation, and truth maintenance techniques. It supports interactive hypothetical model-based reasoning with true dependency backtracking on failure.

In collaboration with the SFU Design Research Group, we are developing a new intelligent design system. Called intelCAD, it combines Echidna with a modern CAD drawing environment and advanced visualization graphics. Two applications of intelCAD are presently under development: an architectural design system (in collaboration with Panabode and Microlight corporations) and a demonstration mechanical engineering system for Sterling engine design (with Dr. John Jones).

Philosophical Issues of Cognition

Dr. Ray Jennings maintains a long-standing interest in the application of weak modal logics to the control of inferences from inconsistent data. Current research uses models for natural conditional logics as formal vehicles for the representation of domain understanding within the knowledge base of a diagnostic system. Such models permit the recognition of any physical relation between objects as inferentially relevant, while clarifying the connection between causal and inferential relations.

Recent work includes the development of an expert diagnostic system language called GALEN based upon higher-order many-sorted logic. GALEN is currently in use at a Veterans Administration hospital in Syracuse, New York, to help refine criteria for EKG diagnosis.

A version of GALEN adapted to industrial applications is being developed by McCarney Technologies in Vancouver, British Columbia, for car engine diagnosis.

Computational Vision

Dr. Ze-Nian Li, Assistant Professor, Computing Science, is interested in parallel and hierarchical computer vision algorithms with massively parallel machines and is also investigating the integration of multisensory data (stereo, range, etc.) for robot vision. Work is also proceeding on the development of knowledge representation and uncertainty management techniques for intelligent systems.

Computational Linguistics

In the area of computational linguistics, Dr. Fred Popowich, Assistant Professor, Computing Science, conducts research focused on unification-based grammar formalisms — highly declarative frameworks that allow the interaction of different forms of information (e.g., syntactic, semantic). Research in this area centers around Head-driven Phrase Structure Grammar (HPSG) and Tree Unification Grammar (TUG).

Particular interest is maintained in the treatment of pronouns within these formalisms and in the problems associated with parsing. Grammar development systems for HPSG and for TUG are currently being developed in projects associated with the Natural Language Laboratory and the Logic Programming Laboratory.

Another area of active research is related to logic grammars, particularly Static Discontinuity Grammars (SDGs). These grammars allow nonlocal relationships to be stated succinctly in terms of sets of definite clause grammar rules. Current work includes the examination of different parsing techniques for SDGs, and various applications of SDGs. One involves the use of SDGs to capture linguistic constraints in an implementation of a Government and Binding Theory grammar for a fragment of a natural language.

For further information, contact:

Dr. N. (Nick) Cercone
Director, Centre for Systems Science
Simon Fraser University
Burnaby, British Columbia
V5A 1S6
Telephone: (604) 291-3229
Fax: (604) 291-4951
NICK@CS.SFU.CA

Research in Biological and Computational Vision at the University of Toronto[1]

By John K. Tsotsos[2]

Overview

Research in biological and computational vision at the University of Toronto (U. of T.) is an interdisciplinary activity, spanning several departments and affiliated with activities at York University. The first instance of computational vision research at the University of Toronto was the 1973 M.Sc. thesis of Nicholas Roussopoulos titled "Abstracting Descriptions from Simple Line Drawings." This work attempted to heuristically decompose closed line curve drawings of simple shapes into meaningful segments, following up on previous work by Guzman. In 1975, Norman Badler completed a Ph.D. thesis titled "Temporal Scene Analysis: Conceptual Descriptions of Object Movements." This is considered the first serious work on high-level motion understanding. The 1980 Ph.D. thesis by John Tsotsos titled "A Framework for Visual Motion Understanding" continued the research begun by Badler and provided both a "big picture" framework for motion understanding systems, as well as demonstrating the concepts with a successful implementation in a rich dynamic domain. These theses were supervised by John Mylopoulos, who also is the "father" of the Artificial Intelligence Group at the university.

A formal vision research group in the Department of Computer Science was begun by Tsotsos in 1980, but biological vision research was ongoing for many years in other departments within the university, primarily in Psychology and Physiology. In 1983, serious cooperation developed between the departments and the informal group RBCV emerged. At that time, the faculty members involved were A.D. Jepson and J.K. Tsotsos, both from Computer Science, P.E. Hallett from Physiology, and P. Kolers from Psychology. There were approximately six graduate students in the Department of Computer Science working on vision. Since then, the group has grown significantly in number and in research activity. Additional faculty are H. Kwan (Physiology), G. Moraglia (Psychology), R. Reiter and E. Milios (Computer Science), G. Hinton (Computer Science and Psychology), and most recently, D.M. Regan (Ophthalmology at the University of Toronto and Psychology at York University). Recent graduates R. Gershon and M. Jenkin, who have remained in Toronto, are also research members. Sadly, Paul Kolers passed away in 1986.

A technical report series was launched in the spring of 1984 and at this writing 20 reports have been produced, involving authors and coauthors from all three departments; this reflects the level of cooperative interdisciplinary work. The technical reports are primarily preprints of papers that have later appeared as refereed conference or journal publications. The graduate student population is about 15 now, primarily students in Computer Science, but also including students from several other departments. Several graduate students have been and are cosupervised across departments.

The group organizes and participates in a number of activities. Peter Hallett organizes the University Lectures in Vision, held yearly and made possible in part by a grant from CAE Electronics Ltd., Montreal. The general theme is to organize a day of lectures and activities around the research interest of a key invitee in biological and computational vision. The 1986 invitee was J.J. Kulikowski, of the University of Manchester, focusing on "Early Form and Motion;" the 1987 invitee was A. Burgess, of the University of British Columbia, with the focus being "Noise and Disorder in Vision." In 1988, the invitee will be F.M. DeMonasterio, who will focus on "Colour Vision." The group also organizes the vision seminars within the Department of Computer Science and runs a bi-weekly vision discussion series.

The research philosophy that guides the group is a simple one: to achieve deep and meaningful results, one must take seriously the hypothesis that human visual information processing can indeed be modeled computationally. This means that: 1) all appropriate biological observations should be carefully considered; 2) all appropriate computational tools are brought to bear on the problem and not just classical mathematics; 3) that both macro-theories ("big picture") and micro-theories

---

[1] Reprinted with permission from *Canadian Artificial Intelligence*, Jan.: 17–23; 1988.

[2] John K. Tsotsos is a professor in the Department of Computer Science with a cross-appointment in Medicine at the University of Toronto. He is a also a fellow of the Canadian Institute for Advanced Research.

(detailed analysis is of individual components) are pursued in parallel, in a complementary and cooperative manner; and 4) that theories produced must not only model the known observations but also have a predictive impact on the experimentation of biological scientists, thus closing the biological motivation loop. A secondary activity of the group is that of experimental computer science, i.e., implementing theories and testing them appropriately, and application of vision results to practical problems.

Research Members

Research members are: Ron Gershon, Expert Systems Center, Allied-Signal Canada; Peter E. Hallett, Physiology, U. of T; Geoffrey E. Hinton, Comp. Sc. and Psychology, U. of T.; Michael R.M. Jenkin, Computer Science, York U.; Allan D. Jepson, Comp. Sc., U. of T.; Hon Kwan, Physiology, U. of T.; Evangelos Milios, Comp. Sc., U. of T.; Giampaolo Moraglia, Psychology, U. of T.; David M. Regan, Ophthalmology at U. of T., and Psychology at York U.; Raymond Reiter, Comp. Sc., U. of T.; John K. Tsotsos, Comp. Sc., U. of T.

Research Projects

Early Motion Sensing and Measurement

N. Lobo (Ph.D. student), J. Tsotsos

In order for early motion measurements to be useful they have to be quite accurate. The computations that recover structure and 3-D motion are highly sensitive to noise. Gradient techniques and spatio-temporal filter methods for computing 2-D velocity degrade when motion occurs over high-contrast textured backgrounds, when there is divergence in the velocity field, and when illumination effects cause the luminance to vary for reasons other than the local translation of a patch. Token-matching schemes may overcome some of these problems but suffer from the perplexing dilemma of choosing an appropriate token. Some of these problems will be studied.

Intermediate Representations for Images

D. Fleet (Ph.D. student), A. Jepson

Previous work concerned the early extraction of visual information. The concern was that richer image information should be extracted without the need for early interpretation or overly restrictive assumptions, such as those inherent in (then) current techniques for motion extraction. This work evolved from a closer look at spatio-temporal properties of center-surround mechanisms in retinal processing. Attention then focused on the construction of efficient (and simple) filters for the extraction of orientation and velocity information based on Fourier analysis and information theory. The work continues now with the construction of families of filters for rich image encodings and the consequences of various hypothetical constraints on such representations. Current interests also include segmentation and shape representation. Questions include what constitutes good form and how such a concept might facilitate figure/ground distinctions and efficient shape encodings.

Models of Image Structure and Primitive Features

M. Langer (M.Sc. student), A. Jepson

The first step in any vision system is the gathering of sensory data in raw pixel form. The fact that the world is highly structured implies that images of the world will also be structured. Any analysis of an image should take advantage of this structure. In particular, since the second step in a vision system is the coding of the image in terms of a set of primitives, these primitives should be chosen with the image structure in mind. One problem then is to form and test a model of image structure. Once this is done, one is able to discuss the efficiency of different image coding schemes with respect to the model. Then, by imposing heuristical constraints that a coding scheme should satisfy, it is possible to compare the efficiency of different schemes. The expectation is that for commonly accepted constraints, such as spatially localized and scale invariant encoding, the most efficient coding schemes will be similar to the ones found in biological systems.

Early Chromatic Processing

A. Jepson, P. Hallett

An information theoretic approach is used to study the functional organization of chromatic processing up to the LGN. There are implications for the choice of spectral sensitivities of the sensors, for algorithms for colour constancy, and for the utilization of simple opponent chromatic units. The eventual goal is to obtain an understanding of the design of the chromatic pathways in terms of a few organizational principles.

Visual Stereoscopic Computation

M. Jenkin, J. Tsotsos, A. Jepson

When we look out on the world, our eyes see views separated by only a few centimetres. This slight difference, and slight differences in ocular orientation, encode the three-dimensional position of objects in the environment. Models for resolving these slight differences (stereopsis) have been proposed by computer vision and psychology researchers. Current models lack both the robustness and sophistication of biological systems. In this thesis, the task of stereopsis will be investigated. The researchers argue against the main tenements of classical stereopsis algorithms, proposing that stereopsis is a spatio-temporal process. They argue against the extraction of complex monocular features such as zero-crossings, against the use of inhibitory surface constraints, and against the use of global stereopsis to produce a unique retinotopic depth map.

A new method for measuring structure from binocular presentations is developed that relates disparity to the local phase difference from bandpass versions of the input images. This disparity measurement produces a dense, robust measurement of structure in space without performing complex token extraction or correspondence. The spatial and temporal properties of the disparity measurement technique are investigated and extensions to the detectors are considered to allow for the construction of detectors tuned to surfaces with particular three-dimensional orientations and to specific three-dimensional trajectories.

Finally, the researchers investigate how the raw disparity measurements should be integrated into a description of the environment. They identify two classes of false targets and propose different techniques for identifying and dealing with them. Rather than construct a unique retinotopic depth map, they argue for the use of an object-based representation and show that even a very simple model-based representation can be used to simplify the global stereopsis problem.

Extracting Egomotion and Environmental Layout from Image Sequences

J. Barron (Ph.D. student), A. Jepson, J. Tsotsos

This thesis concerns the computation of motion and structure from noisy image velocity or from noisy local image velocity information (the Taylor series coefficients of the image velocity field). The two main algorithms developed reconstruct the observer motion and environmental layout (3-D depth information) from time-varying monocular and binocular image sequences. The two algorithms are "unified:" the binocular algorithm reduces to the monocular algorithm if the spatio-temporal baselines are unknown and the left and right image sequences coincide. These algorithms exhibit some novel features — for example, the use of temporal information and the removal of the need to compute point-to-point correspondence in individual image sequences or between stereo image pairs. An extensive error analysis investigates the sensitivity of the motion and structure calculation for various types of input error that includes best, random, and worst case error as well as error in the image velocity means and differences. The researchers demonstrate that increasing the spatio-temporal distribution of image velocities using a least-squares formulation of the algorithms and scaling the Jacobian matrix (among other things) reduces the error amplification. Other results include the presence of dual solutions, the presence of multiple (nondual) solutions, the effect of violating the algorithms' underlying assumptions, the relationship between error in image velocities and their corresponding Taylor series coefficients, and the effect of using normal (versus full) image velocity fields. In the binocular case the researchers also investigate the effect of varying the spatio-temporal baseline and of parallel versus convergent-divergent stereo setups.

Self-organizing Feature Maps

A. Jepson. M. Langer (M.Sc. student), D. Fleet (Ph.D. student)

This is an integrative project extending and unifying the results of the previously listed projects. The goal is to obtain realizable unsupervised learning rules for image feature maps based on information theoretic principles.

The REx Robot Rover Project

E. Milios (project leader)

The goal of this project is to design an image understanding system in which to integrate much of the research being conducted by the University of Toronto vision group. The problem area is the task of navigating a camera-equipped robot successfully in an office environment containing moving people and other objects. Low-level vision tasks will be performed by a selection of independent modules that compute various retinotopic (image-coordinate-based) "feature maps" from the sequence of input images. These maps form the input to the high-level portion of the system, in which recognition of instances of known objects occur.

Much of the research required is in the design of representation and reasoning mechanisms that extend the state of the art in high-level vision. In particular, it is necessary to address the problem of identifying partially occluded objects and to increase the generality of object representations to model motion characteristics and resulting shape and position changes in a general manner. The high level will incorporate attention mechanisms to direct the operation of the low-level modules (which modules to use, resolution at which to compute, area of image).

As the system matures, it is expected that the work of additional department members may be integrated. Experimentation will be facilitated by modularity and careful documentation. Greg Dudek, Mike Jenkin, Howie Marcus, Evangelos Milios, and David Wilkes are currently building a working prototype that addresses all stages of the computation but uses range data instead of visual data. After a better understanding of the overall system design issues is achieved, range data will be replaced by visual data.

Process-based Shape Theories

E. Milios

Milios is currently working on computational implementations of process-based shape theories suitable for describing growth or evolution of amorphous shapes. His research is motivated by the recent emergence of psychological theories that view shape as an outcome of deformation processes that introduce curvature extrema and inflection points and by the fact that current shape matching methods either work on polygonal/polyhedral approximations or search for exact matching between smoothly curved shapes. His approach is based on higher-order polynomial shape modeling and segmentation and attempts to match segments at multiple abstraction levels using curvature-based representations. Shape matching is viewed not only as establishing correspondences between segments of two shapes, but also as the explanation of their differences by means of curvature processes. This research has potential application to the automatic interpretation of spatio-temporal satellite imagery for weather forecasting and to the description of growth of biological phenomena.

The Logic of Depiction

R. Reiter, A.K. Mackworth (University of British Columbia)

The researchers propose a theory of depiction and interpretation that formalizes image domain knowledge, scene domain knowledge, and the depiction mapping between the image and scene domains. This theory is illustrated by specifying some general knowledge about maps, geographic objects and their depiction relationships in first-order logic with equality.

An interpretation of an image is defined to be a logical model of the general knowledge and a description of that image. For the simple map world the researchers show how the task level specification may be refined to a provably correct implementation by invoking model preserving transformations on the logical representation. In addition, they sketch logical treatments for querying an image incorporating contingent scene knowledge into the interpretation process, occlusion, ambiguous image descriptions, and composition.

This approach provides a formal framework for analyzing existing systems such as Mapsee and for understanding the use of constraint satisfaction techniques. It can also be used to design and implement vision and graphics systems that are correct with respect to the task and algorithm levels.

Finite Automata Models for Cortical Texture

M. Hofmann (Ph.D. student in Zoology),
P. Hallett, A. Jepson

The main focus of interest is the area of pattern and texture discrimination and segregation. The research is composed of two parts. A psychophysical study using human observers will test the ability of subjects to discriminate patterns. The particular class of patterns used in this study are composed of superimposed sinewave gratings of different orientations. The spatial frequency, orientations, and relative phases of the patterns can be varied to produce a wide range of patterns.

The second part of the research is the construction of a computer simulation based on known physiology of the visual system. The purpose of the model is to try to simulate the visual system's ability to discriminate and segregate textures and patterns. The results of the simulation will be compared with the results from human observers given similar input patterns.

Evaluating Blood Perfusion from Retinal Images of Diabetes Patients

G. Crookall (M.Sc. student), C. Francis (clinical clerk in the Dept. of Medicine), A. Kurusek (Dept. of Ophthalmology, Toronto General Hospital), J. Tsotsos

This project will develop a vision system to quantify and identify patterns of anomalous blood perfusion from images of the human retina. In diabetes, blood perfusion in the retina is affected, leading to the potential for partial or total blindness in the affected eye. Some observable problems include localized changes in colour; localized changes in image texture of the retina, changes in the pattern of arteries, and changes in the size and elasticity of arteries. Combinations of anomalies have specific meaning as well. The patterns and characteristics of those anomalies and their changes during the time course of the disease are currently topics of active clinical research. It is difficult, however, to quantify texture or size changes by observation alone. Therefore, a research tool is required to assist with the automatic and objective evaluation of retinal images. The types of anomalies will be catalogued and appropriate techniques developed for extracting image features and temporal feature changes and then matching features to anomalies believed to be manifestations of specific disease stages. This research tool will permit clinicians to pose hypotheses about disease progress and effects on the retina and will provide the quantitative analyses necessary to compare hypotheses. The control structure will be based on that used in the ALVEN system developed by Tsotsos for evaluation of the dynamic performance of the human left ventricle from X-ray images.

Hierarchical Time-varying Cooperative Computation

E. Shavit (M.Sc. student), J. Tsotsos

In Tsotsos' 1980 Ph.D. thesis, a time-varying, hierarchical relaxation labeling process was presented. It performed very well in that application. Although an empirical study of the process was conducted (it seemed too difficult to analyze analytically at the time because it was cast in the old-style relaxation formalism) and many properties were demonstrated, such as convergence and compatibility settings related to the semantics of hierarchy organization, there were several issues remaining. The foremost issue is that of characterizing the computation, i.e., what exactly is this large system computing? The current project will attempt to reformulate the computation as an optimization problem and address this question from an energy minimization point of view. Note that all previous relaxation algorithms make a critical assumption that does not hold in a time-varying domain — namely, that an initial solution is available as a starting point that is "close" to the desired solution. In a domain whose input is time-varying and where the structure over which the computation is being performed is also time-varying, this assumption is violated much of the time. The previous formulation dealt with this problem by using an executive controller to continually tune and adjust the system's parameters. It is expected that such a solution will still be necessary; moreover it is desirable because other research (Tsotsos' complexity level analysis of vision) has shown that exactly such an "attentive" component is a necessary one for human vision.

The Use of Curvature in Model-based Vision

G. Dudek (Ph.D. student), J. Tsotsos

The use of surface shape as an aid to object recognition is being studied. The particular measurable quantity proposed for use as a shape indicator is local surface curvature. It has been demonstrated that in a constrained environment local surface curvature estimates can be obtained from surface shading gradations. Unfortunately, the accuracy of the recovered parameters is inadequate for the true surface shape except in the most elementary conditions. Non-Lambertian surfaces or nearby light sources are among the factors that can lead to corruption of the results. The goal of the present research is to use local curvature estimates to classify patches of the surface using a quadric surface model. Furthermore, by carrying out this classification at the multiple scales at which curvature changes manifest themselves, different types of co-occurring surface structure should manifest themselves explicitly. The curvature information is explicitly available in the gradient of the difference of Gaussians of the image. Using this inherently multi-scale representation, the surface may be classified on a local basis. Once this local classification is performed, the measurements must be organized to describe the surface in terms of patches of varying size. It is hoped that topological laws can be used to aid in this process.

Finally, these surface patches at multiple scales are to be linked together into a coherent representation of the object in 2½-D. Although this representation will not fully capture surface shape, it should capture the major structural features. An object recognition process based on this representation is now being planned. It should embody aspects of the Grimson/Lozano-Perez style constrained matching algorithm at each scale, while capitalizing on the multi-scale nature of the representation by using a coarse-to-fine matching paradigm.

Associative Computation in Vision

D. Wilkes (Ph.D. student), J. Tsotsos

The researchers are interested in the development of parallel architectures for vision. Work in parallel distributed implementations of various filtering, association, and constraint satisfaction tasks suggests the plausibility of a complete vision architecture based on such implementations and their extensions. Important ideas to be incorporated in such an architecture are the use of attention, 2-D perceptual cues, and representation of objects in terms of sequences (or, more generally, graphs) of cue areas separated by arcs representing the displacement from one cue area to the next. It is desired to achieve reliable initial object recognition without resorting to depth or surface orientation reconstruction. The proposed architecture draws on the vision models of Triesman, Tsotsos and Lowe. Parallel computation of "feature maps," using simple filtering operations, provides cues for initial bottom-up direction of attention to cue areas in the image data. Model elements detected in the cue areas influence direction of attention by prompting, e.g., edge-following jumps of attention. Cues also prompt associative retrieval of object models, which then provide a top-down influence on attention. Image to object-based coordinate transformation is determined concurrently by matching of cue pairs in the image and object model to each other. Problems to be solved are primarily concerned with object representation. Reliable cues must be combined in a way that allows the expression of 3-D spatial information and associative model retrieval based on a small number of cues with their relative image positions. Extensive use of statistical inference, based on prior experience of the system, will be incorporated.

Analyzing Vision at the Complexity Level

J. Tsotsos

"Grand theories" of perception are easy prey for criticism, criticism that in one important sense is unfair at this point in the development of this discipline. There is no test that can be applied to a grand theory in order to determine whether or not fundamental considerations are satisfied. The researcher's work proposes that satisfaction of the space and time complexity constraints be one of the elements of the test that new theories of visual perception must pass. This ongoing research attempts to tie many biological observations of human and primate vision systems with the thread of complexity satisfaction, a novel approach to this problem. The goal is to develop some of the key design principles for visual systems, to discover the constraints on visual system architectures that permit space and time complexity satisfaction, and to produce a theory with predictive impact, both anatomic and psychological. Much of the motivation for this research comes from the psychology community. Neisser, for example, claimed that a visual theory based solely on spatial parallelism was doomed to failure because the brain was not large enough. This is why attention was required. Although true, this statement does not lead to any useful constraints on visual system architectures, but it hints at the important complexity issues that to this point have not been adequately addressed.

Computational complexity issues are broad and pervasive in the development of a theory of perception based on the premise that human visual perception can be modeled computationally.

Since complexity theory is a major contribution of computer science, its use is natural within the computational paradigm. It is further hypothesized that evolution, where guided by basic design principles, tended to prefer system designs with lower computational complexity. If these hypotheses are true, then complexity analysis should reveal basic insights into the structure and performance of human vision. The serious consideration of computational complexity is claimed to be critical in the computational modeling of perception, and indeed, in the computational modeling of any aspect of intelligence. One of the key problems with AI is that the solutions proposed are so fragile with respect to "scaling up" with problem size. In particular, parallel solutions, such as those proposed by the connectionist community, although motivated by complexity considerations, typically fail to demonstrate the computational sufficiency of their approaches. It is also interesting to note that the application of complexity analysis in order to uncover basic limits of the information processing capacity of the visual system did not appear at all in the overview of a large number of techniques and approaches to discovering the limits of perception in *Limits in Perception* (edited by van Doorn, van de Grind, and Koenderink, 1984).

There are three major results that are derived using space and time complexity analysis, a technique that may be one of the most important contributions of computer science with theoretical implications for a wide range of sciences. First, a set of architectural constraints is presented that satisfies the space complexity of the human visual system and the time complexity of preattentive vision. These constraints are derived using two guiding principles claimed to be critical in the development of any theory of perception. The first is that analysis at the "complexity level" is necessary to ensure that the basic space and performance measures observed in human vision are satisfied by a proposed system architecture. Then, the "maximum power/minimum cost principle" ranks the many architectures that satisfy the complexity level analysis and allows the choice of the best one. The best architecture chosen using this principle is completely compatible with the known architecture of the human visual system, and in addition, leads to several anatomical predictions. Secondly, within this architecture, a new explanation for visual search performance is proposed that differs in important ways from Treisman's account and is accompanied by several predictions. The key result regarding this explanation is that preattentive vision is shown to be just a special case of visual information processing and not a processing stage distinct from attentive vision. Thirdly, the inherent computational limits of this bottom-up architecture are exposed and thus provide a computational justification for the necessity of attentive processing in visual theories, something that has been debated in the computer vision community and is counter to Marrs' view of vision.

Visual Search Performance in Humans

R. Zemel (M.Sc. student), J. Tsotsos, G. Moraglia

There are a number of key implications on visual search performance in Tsotsos' complexity level theory. A number of experiments may be devised to test each prediction, and the psychology experimentation is being done under the guidance of Moraglia. Possible hypotheses include: map activation is not the only mechanism active for pop-out displays; an increase in the number of active maps leads to an increase in response time for pop-out displays; strength of response can order candidates for the match-to-target process; non-distractor elements in a conjunction display have no effect; serial search is required to identify and count more than one target in a pop-out display; no spatial relations pop-out; as the number of targets increases, so does the response time.

Learning through Amplification of Correlations and Fluctuations in a Modular and Hierarchical Neuronal Network

T. Yeap (Ph.D. student in Electrical Engineering), H. Kwan, J. Tsotsos, S. Zaky (Electrical Engineering)

Recent progress in the study of brain and VLSI technology has led to growing interest in research on massively parallel architectures. These machines use many simple processing elements (PE) connected to each other. A weight is assigned to each connection, and the machine is programmed by adjusting these weights. Two well-known examples are the Boltzmann machine and Backward Error Propagation Machine, and much work has been done on learning algorithms for such networks. However, an implicit assumption in past work has been that an initial value for the recognition or learning task is available that is close to the equilibrium state of the network. This is a useful assumption, valid in many tasks; however, in recognition of a time-varying domain, it is often violated (in vision, for example, when a moving object stops). This thesis examines the possibility of learning and recognition in a network that may operate both near to and far from equilibrium. It shares many characteristics with Boltzmann architectures; differently, however, the PE considered is an analog computation unit. It has an analog integrator, a threshold unit, and a differential positive feedback loop. As a result, the circuit model is sensitive to change or fluctuation in the network. With the ability to amplify fluctuation, the whole network is very unstable, and thus it is very easy for the system to go into a state of chaos, which is a common process far from equilibrium. It has been hypothesized that nature achieved order through the use of a competition process between neurons, and such a competitive mechanism is included in this design.

The learning algorithm is based on the study of self-organization in nonequilibrium systems done by Prigogine. According to Prigogine, most natural systems that are capable of self-organization have the property of operating far from equilibrium. Order is achieved through the ability of the system to amplify fluctuations and correlations. Learning in a network is studied by incorporating fluctuation, correlation, and competition as a process of adjusting the weights of the connections between PEs. Several theories on self-organization in nonequilibrium systems are currently being studied.

For a large network performing a complex task, long computation time is required to settle to a unique solution. To reduce the computation time, task partitioning is crucial. Units may thus be organized hierarchically.

Auditory and Visual Psychophysics and Electrophysiology

D.M. Regan

Regan is involved in a number of research areas: 1) experimental visual psychophysics and mathematical modeling for form from motion, stereopsis, motion in depth, and colour — a well-equipped vision laboratory including eye-movement recording facilities is available; 2) auditory psychophysics and modeling: FM/AM channels and neural processing of speech — a sound-insulated room is available; 3) visual, auditory, and inter-modality magnetic and electrical recording from the human brain, using a 7-channel SQUID neuromagnetometer in a magnetically shielded room; 4) studies of visual and auditory disorders via psychophysical and evoked brain potential and neuromagnetic investigations; 5) applications of psychophysics to hand-eye coordination in aviation and space with opportunities to work with flight simulators and high-performance jet aircraft.

Preattentive and Attentive Vision

G. Moraglia

Current research focuses on the issues related to the preattentive–attentive dichotomy in human spatial vision. This dichotomy is explored within the context of experimental tasks involving search, detection and recognition of visual signals in noisy backgrounds. The aim of this research is to contribute to the development of alternative theoretical accounts of this dichotomy. Allied concerns include the modeling of human performance in close experimental analogs of real-life pattern recognition tasks.

Visual Information Processing

G.E. Hinton

Hinton has worked on connectionist models of viewpoint-invariant shape recognition and on general models of relaxation search in massively parallel networks. He is currently interested in the issue of what frames of reference are used for human vision and motor control and is trying to devise efficient connectionist methods of performing coordinate transformations between different reference frames (e.g., between retinocentric and object-based reference frames.) He is also interested in the application of unsupervised connectionist learning procedures to the task of constructing a hierarchy of representations in early vision.

Research Facilities

There are a number of laboratories to which any member of the group has access, each providing a unique set of research facilities. These include:

- several Symbolics LISP Machines
- several colour displays with attached array processors

- several Sun systems
- image digitization facilities
- DEC VAX 11/780
- single-cell recording facilities for primates
- computer-driven 4-field monocular tachistoscope
- computer-driven 2-field binocular tachistoscope
- 7-channel SQUID neuromagnetometer with shielded room
- eye movement recording equipment

Research in Biological and Computational Vision Publication List

The reports are available upon request from: RBCV Technical Report, Secretary, Department of Computer Science, 10 King's College Road, University of Toronto, Toronto, Ontario, M5S 1A4. Several of the abstracts have been published in the October 1987 issue of this magazine.

Section 5

AI in Business and Industry

Acquired Intelligence Inc

Keywords: knowledge acquisition, expert systems, intelligent tutoring systems, machine learning systems, pattern recognition, preventive debugging, structured induction, top-down problem decomposition.

Acquired Intelligence Inc (AI Inc) is a privately held, Canadian-owned corporation involved in the research, development, and commercialization of artificial intelligence technology. AI Inc offers products and services in the area of expert systems and intelligent tutoring systems. This article describes AI Inc's development and exploitation of AI technology as well as ongoing research and development at the company.

Central to the application of expert systems and intelligent tutoring systems is the issue of *knowledge acquisition*. A long-term research project at AI Inc has resulted in a major breakthrough in knowledge acquisition technology. This breakthrough forms the basis of the company's products and provides a unique benefit to its clients. Continuing R&D efforts ensure future product advances and service opportunities.

Products and Services

Tools

AI Inc produces expert system tools that can be used directly by experts to build their own expert system applications. ACQUIRE™ is an expert system development environment that includes a knowledge acquisition module with machine learning capabilities, an expert system shell, and a variety of utilities for designing and producing various types of reports. This system, first and foremost, addresses the difficulties inherent in acquiring knowledge from experts, and is designed to enable experts to develop and maintain their own knowledge bases on personal computers. It is easy for experts to use because its representation scheme is based on fundamental knowledge acquisition considerations. In ACQUIRE, knowledge acquisition issues determined the nature of the knowledge representation rather than the representation determining the acquisition strategy.

The system uses a structured approach to guide the expert through the knowledge elicitation process, capitalizing on natural pattern recognition skills to facilitate the acquisition of knowledge. Step-by-step, knowledge is decomposed into subtasks that are manageable and meaningful to the expert. The integrity of expertise is maintained under the expert's, not a programmer's, control. The "knowledge-acquisition-friendly" representation scheme produces extremely thorough and complete knowledge bases, a first step toward preventive debugging. The system also incorporates machine learning capabilities, based on the latest advances in structured induction, to improve the organization of knowledge through suggestions to the expert. A sophisticated expert system shell optimized for processing cases through ACQUIRE knowledge bases, and a set of utilities for building databases of cases and designing structured reports, are also provided. A runtime version is available for executing completed applications.

With ACQUIRE, help from knowledge engineers or special training in expert systems technology is not required. The system, then, should contribute to the widespread proliferation of knowledge-based systems technology in any organization. The system has already been thoroughly tested in a wide variety of applications in forestry, mining, market analysis, medicine, space, agriculture, and neuropsychology. The system is implemented in C and runs in MS-DOS or Unix operating systems on IBM PC computers. Conversions to other platforms is underway.

AI Inc also distributes the LISP intelligent tutoring system developed at Carnegie-Mellon University.

Applications

AI Inc develops expert system applications as solutions to identified problems both as customized applications and as applications that are marketed to targeted business and professional communities. Examples of the latter include expert systems for neuropsychology, school psychology, haematology, and market analysis.

Services

AI Inc provides a wide range of services related to knowledge acquisition, expert systems, machine learning, and intelligent tutoring systems technology including: consulting, surveys, educational services, feasibility studies, and development of prototype and commercial systems. Service contracts are selected to augment product capabilities through technical enhancements and also to increase the market potential for the company's products through the addition of other product capabilities.

Examples of applications areas are: dairy operation; safety analysis; market analysis; combat intelligence; veterinary diagnosis; haematology screening; educational screening; mineral exploration, geophysics, and geology; plant equipment operation and failure analysis; and neuropsychology, brain dysfunction, and learning disabilities.

Research and Development

Knowledge Acquisition and Machine Learning Integration

ACQUIRE currently uses a subset of AI Inc's C-DUCE machine learning system to augment aspects of the knowledge acquisition process. The capabilities of the system are being extended to make the structured induction approach more amenable to financial and engineering problems where number handling capabilities are critical. The enhanced version of the machine learning system, to be called ACQUIRE-MLS, will operate in conjunction with the ACQUIRE expert system shell.

The enhanced machine learning system is also being integrated with the knowledge acquisition module. This will allow a database of examples to be used for the initial construction of a knowledge base using machine learning techniques, and allow further elaboration, refinement, and reorganization by the expert using ACQUIRE. In this environment, knowledge bases will be a composite of both expert- and machine-generated knowledge, extending expert systems technology to problems that cannot be handled adequately by either approach alone. The integrated system is being field tested at MacMillan Bloedel Ltd.

Knowledge Acquisition From Text

The acquisition of knowledge from written materials is routinely performed by people who consult regulations, handbooks, and reference materials in many areas of employment. AI Inc and faculty from the University of Victoria's Linguistics Department are involved in a research project to develop a system that can construct ACQUIRE knowledge bases from analysis of written materials. This project has focused on the acquisition of knowledge from books of regulations such as those produced by government departments and agencies. To date, the prototype has been used to derive a knowledge base from an analysis of City of Victoria parking by-laws. This type of system could be an enormous aid for monitoring compliance, providing guidance in the interpretation, improving the consistency, and facilitating the delivery of regulations (e.g., in the construction and transportation industries).

ACQUIRE Intelligent Tutoring System

AI Inc has been awarded a contract from the Canadian Space Agency's Strategic Technologies for Enhanced Automation and Robotics (STEAR) program to build a knowledge-based intelligent tutoring system based on the ACQUIRE system. ACQUIRE-ITS will enable computer-naive users to develop intelligent training applications in their areas of expertise. This environment will be capable of being tailored to meet both the crew training demands of the space station's Mobile Service System and the training needs of Canadian industry. AI Inc believes that widespread proliferation of knowledge-based technology is possible with the emergence of cost-effective tools that run on affordable computers and enable experts to develop and maintain their own applications. Therefore, the intelligent training applications are being developed directly by experts on affordable computers.

A *Curriculum Author* module enables experts to develop a curriculum composed of lessons in a case-study approach. The system is oriented toward "learning by doing." Students learn by solving actual cases which have been packaged into lessons by the expert to demonstrate important and related concepts. By inspecting the structure and content of the expert system's knowledge base the Curriculum Author can help the expert select the best cases to package together into lessons.

A *Tutoring Module* monitors a student's progress in solving cases and constructs a representation of the student's growing knowledge as a "naive" knowledge base. The Tutoring Module gradually matches the student's naive knowledge base with the expert's by identifying incorrect rules, gaps in the knowledge, and, most importantly, poor structuring of the knowledge (i.e., incorrect packaging of rules), which indicate misconceptions about how knowledge should be used. The Tutoring Module accomplishes this merger by selecting and presenting appropriate cases, from the lessons constructed by the expert, to help the student elucidate the correct knowledge in active problem solving. Otherwise, the Tutoring Module explicitly demonstrates differences between the student's knowledge base and the expert's.

For further information contact:

Dr. Beverly J. Smith
Acquired Intelligence Inc
#205-1095 McKenzie Avenue
Victoria, British Columbia
V8P 2L5
Telephone: (604) 479-8646
Fax: (604) 479-0764

Applied AI Systems, Inc

Description of Company

Applied AI Systems, Inc (AAI) provides expertise in the research, design, implementation, and deployment of practical AI systems in the fields of government operations, engineering, transportation, communication, and business.

AAI educates and trains trainees in international AI business and in the use and practical application of AI tools. The company also provides technical and business consultation to the client for a specific period of time or on a specific theme. AAI is a distributor for several selected AI system development tools.

Description of Some Projects at AAI

Expert Systems in Canadian Transportation

With COGNOS, conducted investigation of how expert system technology could be applied to the Canadian Transportation sector. Report published, *Expert Systems: Their Application in the Transportation Sector.*

Voice Interactive Technology

For Transport Canada, surveyed voice and natural language technologies as related to transportation applications. Report, *State of the Art of Voice Interactive Technology and its Application to the Canadian Transportation Sector.*

AI in Marine Operation

Detailed case study for the automation of navigation, fuel-efficient operation, and on-board in-operation maintenance. Published Report, *A Feasibility Study of Expert Advisors to the Operation of CN Marine's M.V. Caribou.*

AI Research in Japan, Britain, and Europe

Reports published reporting visits to applicative AI R&D centers.

Sixth Generation Computing

Papers presented at the 2nd International Symposium on Artificial Intelligence and Expert Systems in Berlin and Conference on the Impact of Artificial Intelligence on Business and Industry in Texas.

Advanced Natural Language Research

Review of natural language research and multimodal intelligent interface research in the United States.

Natural Language Understanding in Neural Network

Survey of research activities in neural-network-based natural language methods.

Fusion of Symbolic and Neural AI Paradigms

Approaches for interfacing symbolic and neural AI paradigms. Survey, design, and implementation paper presentation at AAAI 90.

AI System Applications within Government

Paper presented at IEEE Third Expert Systems in Government Conference, 1987.

Development and Demonstration of Knowledge-based Vehicle Dispatcher's Assistant

Prototype which concentrated on development of effective multimodal user interface in transport vehicle dispatching. Reported at 1989 Canadian Conference on Electrical and Computer Engineering and the session on electronic Revolution in Goods Distribution Systems at TRB Conference in Washington.

Demonstration of a Contract Evaluation Expert System

Prototype knowledge-based system to assist less experienced contracting officials apply a uniform set of rules when awarding contracts.

Expert System Demonstration in Urban Transit

Demonstrated the applicability of knowledge-based systems in spare-board management in urban transit. Paper presented at IEEE AI 88 in Japan and the Canadian Conference on Electrical and Computer Engineering in Vancouver.

Intelligent Text Processing

Report on the short- and long-term trends in online natural language technology and an evaluation of products currently on the market.

Neural Network Applications

Report on the broad implications of neural network technology and the scope of future neural network applications.

Autonomous Vehicles

Prototype completed for demonstration of a real-time manoeuvering system for autonomous vehicles. Paper presented at IEEE International Robots and Systems Conference (IROS 90).

Intelligent Toys

AAI is researching and making recommendations for the implementation of advanced AI techniques in the toy market of the future jointly with the Massachusetts Institute of Technology's Intelligent Insect research group.

Fleet Operations Resource Management Decision Advisor

With Fleet Technology Limited, Phase II is approved for an Advisor for the Department of Fisheries and Oceans in making decisions for Capital and O&M programs for the most efficient use of ships.

For more information, please contact:

Ann Griffith
Applied AI Systems, Inc
Gateway Business Park
340 March Road, Suite 500
Kanata, Ontario
Telephone: (613) 592-3030
Fax: (613) 592-2333

Bell-Northern Research[1]

By Dick Peacocke[2]

Bell-Northern Research (BNR) is the research and development organization for Northern Telecom and Bell Canada, with more than 6700 staff in laboratories across Canada, the United States, and in the United Kingdom. Speech and image processing have been the focus of research for more than a decade, and the artificial intelligence program has been running for four years.

The scope of the AI program is to explore and develop AI techniques for use in products and services provided by BNR and its parent companies and for use in design and development processes. A core AI group was formed at the beginning of 1984 to handle basic technology issues and to help define and support applications throughout the corporation. The core group is also responsible for liaison with AI activities in universities and government. This liaison is viewed as an important task, leading (it is hoped) to collaboration in some instances.

Telecommunications and information systems are major products of Northern Telecom. As they become more and more sophisticated in their capabilities in the future they must still be easy to use and manage. A central theme for AI work at BNR is coupling machines and people together in the performance of complex tasks.

At Bell Northern Research we believe that many of the tasks involved in establishing and operating telecommunications networks can be done using expert systems, and indeed we believe that these types of system will soon be essential in helping people perform the complex network tasks required of them. Examples of such tasks are planning and design, engineering, operations, maintenance, and administration.

Complexity of the telecommunications network is increasing rapidly with the introduction of new functions and services, increased connectivity, and the advent of new components. This means that fewer and fewer tractable algorithms exist for the tasks listed above. Expert systems allow implementation of procedures where more emphasis has to be placed on our experience and what we know about situations and events surrounding them, rather than on formal sequential algorithms. After articulating and transforming our knowledge of the network and its operations into suitable representations we can then apply that knowledge to the tasks that need to be done. Our aim is to enhance the scope and effectiveness of these tasks, increasing the quality of work performed and reducing the effort required. This has an impact on documentation and training, and it will increase the ability to absorb new technology.

The telecommunications network is a particularly fertile area for applying expert systems because the network and its equipment form a highly-structured synthetic domain. (This compares with domains such as medicine or geology where the underlying rules of operation are based on natural phenomena, and extensive — often visual — descriptions of these phenomena are required.) Some knowledge-based prototypes that have been built at BNR are shown in Table 6. Interesting issues include domain characterization, knowledge elicitation and representation, development of methodologies, selection of tools, and prototyping and testing expert systems.

Two of the entries in Table 6 are not really expert systems. They represent work in natural language processing and work in applying AI to software engineering.

Natural language research at BNR is aimed at intelligent human–machine dialogue. Natural language has been studied from both the computational and behavioral points of view. The strategy used is the "entire system strategy," viz. building and using entire natural language programs rather than concentrating on isolated phenomena. An experimental domain-independent natural language interface has been developed for an existing entity-relationship database management system. Both the in-house and a commercial natural language system have been used to study behavior of users interacting with an operational database. Issues of interest include different modes of dialogue, how context is used, and dynamic patterns of focus.

As well as research in expert systems and natural language processing, we have been investigating

---

[1] Reprinted with permission from *Canadian Artificial Intelligence*, October: 32–34; 1988.

[2] Dick Peacocke is Manager of the Knowledge Technology group at BNR. He is President of CSCSI/SCEIO and Adjunct Professor at Carleton University. He has a Ph.D. in computer science from the University of Toronto.

Table 6. Knowledge-based prototypes at BNR.

Application	Product area	Function
Switch maintenance	DMS-1OO™ family of digital switches	Provide advice on trouble-shooting to technicians operating a digital telephone switch
Access maintenance	ISDN access	Provide advice to customers on loop access equipment and procedures
Monitoring and maintenance	DATAPAC™	Collect and analyze alarms from a packet switching network. Recommend appropriate maintenance actions actions
Configuration of equipment	Business Communications Systems	Transform user and site requirements as entered on a form into detailed hardware
Configuration of private data networks	Bell Canada data communications services	Choose the most appropriate data service(s) based on clients' terminals, hosts, protocols, transaction volumes, etc.
Circuit design engineering	Bell Canada Special Service Digital Network	Automate the circuit design process for certain types of special service circuits
Software engineering	DMS-100 software life cycle	Use a knowledge base to capture software design input. Derive documentation, test cases, design checks, etc. from the knowledge base
Natural language interface	Use on relational or entity-relationship databases	Provide a processor for database queries expressed in English instead of in a formal query language like SQL
Screening electronic mail	Business Communications Systems	Monitor, creeen, and take action on mail based on organizational knowledge and personal rules specified by users

techniques and algorithms for automating the design and development of software systems and improving development productivity. This work places particular emphasis on specific BNR domains. The term "development" as used here includes all of the software life-cycle phases from requirements and specification through to testing, release, and maintenance. Topics of interest include rapid prototyping, intelligent development environments, and automatic test generation.

As computing environments, we have been using Xerox LISP machines with InterLISP-D, KEE,™ Prolog, S.l,™ and LOOPS,™ Prolog on a proprietary 68000-based family of workstations, Prolog and LISP on IBM mainframes, and several small shells for PCs and the Macintosh.

References

1. Ali, Y., Aubin, R., and Hall, B., "A Domain-Independent Natural Language Database Interface" *Proceedings 6th Canadian Conference on Artificial Intelligence,* Montreal, Canada, May 1986, pp. 62–66.

2. Aubin, R., Peacocke, R., and Zlatin, D., "From Specifications to Test Cases: AI in Software Quality" *Proceedings IEEE International Conference on Communications,* Toronto, Canada, June 1986, pp. 1176–1179.

3. Bult, T., Peacocke, R., and Snarr Carter, V., "An Interactive Expert System for Switch Maintenance" *Proceedings International Switching Symposium,* Phoenix, Arizona, March 1987, pp. 59–65.

4. Ferguson, I., Rabie, S., Kennedy, J., and Peacocke, R., "A Knowledge-based Sales Assistant for Data Communications Networks" *Proceedings IEEE International Conference on Communications,* Seattle, Washington, June 1987, pp. 1634–1642.

5. Ferguson, I., Zlatin, D., "Generic Strategies and Representations for Communications Networks Sales" *Proceedings 7th Canadian Conference on Artificial Intelligence,* Edmonton, Alberta, June 1988, pp. 249–254.

6. Peacocke, R., "On-Line Expertise for Telecommunications" *Expert Systems for Telecommunications,* Liebowitz, J., (editor), John Wiley, New York, February 1988.

7. Peacocke, R., Rabie, S., "Knowledge-Based Maintenance in Networks" *IEEE Journal on Selected Areas in Communications,* Volume 6, Number 5, June 1988, pp. 813–818.

8. Rabie, S., Bult, T., Snarr Carter, V., Cavan, B. "DRESS: Diagnostic Repair Expert System Shell" *Proceedings International Conference on Communications,* Philadelphia, Pennsylvania, June 1988, pp. 1467–1472.

9. Rabie, S., Rau-Chaplin, A., Shibahara, T., "DAD: A Real-Time Expert System for Monitoring of Data Packet Networks" *IEEE Network Magazine,* September 1988.

CAE Electronics Ltd.

Introduction

CAE's activities in Artificial Intelligence began in 1985 with an evaluation of AI's potential contribution to CAE's needs and interests, particularly in realtime systems. Dr. Renato De Mori of McGill University was consulted to determine the state of the technology and the most effective means of technology transfer. As a result of these initiatives, several projects have been undertaken:

- APACS: An Advanced Process Analysis and Control System. This R&D project is in partnership with Ontario Hydro, Shell Canada, and the University of Toronto and is funded through PRECARN Associates.
- FLITE: The FLight Intelligent Troubleshooting Environment, an R&D project in partnership with CRIM and McGill University.
- ITEMS: An Interactive Tactical Environment Management System which includes a rule-based component.
- ATC Simulation: Knowledge-based techniques for speech recognition in air traffic controler training.

These projects are discussed in greater detail below.

In June 1988 CAE formed a specialized AI group with a mandate to conduct existing projects and pursue new applications. At present, the group consists of three engineers with advanced graduate training in AI theory and technologies and a total of 14 years of experience at CAE in AI, system development, simulation, and control systems. Working relations are maintained with McGill University and the University of Toronto.

APACS

The Advanced Process Analysis and Control System (APACS) project is a five-year program of research and development supported by PRECARN Associates. The participating organizations are CAE Electronics, Ontario Hydro, Shell Canada, and the University of Toronto. A six-month feasibility study, funded by PRECARN, was completed in June 1989. Based on the results of the feasibility study, the APACS project has been approved for funding by PRECARN's Board of Directors.

The Problem Domain

Continuous process control technology has evolved from a series of manually-oriented tasks to a system of sophisticated computer-based operations. Despite impressive progress, further automation of these facilities still poses challenging research and development opportunities. For example, Canadian nuclear power generating units running under normal operating conditions are currently controlled by computers. However, in the event of abnormal operating conditions, human intervention becomes necessary since the existing automatic control algorithms cannot always cope with abnormal situations.

Advances in two technical areas will have a significant impact on the future operation and control of continuous process facilities:

1. Techniques from artificial intelligence could be used to develop knowledge-based systems with the ability to support technical decision-making in a manner similar to human operators.
2. Due to increases in the processing capabilities of computers, the possibility exists for a "faster than realtime" simulation capable of supporting on-line realtime diagnosis of system faults and the prediction of plant behavior.

While the feasibility of applying AI techniques in the power industry has been demonstrated by the success of several fielded and prototype task-specific expert systems, the APACS research program presents many challenges in the plant-wide application of this technology and in the integration of expert systems with conventional plant control systems.

Objectives of the Research Program

The principal objectives of the research program may be summarized as follows:

- To design and develop knowledge-based systems to deal with operation and control problems within a power plant.
- To design and develop an advanced simulation facility, the plant analyzer, to support the realtime diagnosis of plant malfunctions and the prediction of plant behavior.

- To design and develop an operator–machine interface in compliance with the operating requirements of a control room environment.
- To integrate these systems within an advanced process analysis and control system for a power plant.

Significance of the APACS Project

The APACS research program will address several difficult AI research issues. Advances in AI techniques are expected in the following areas:

- knowledge representation for process system structure and behavior
- realtime control of inference
- adaption of expert system building tools for realtime applications
- integration with existing control systems and plant analyzer

An economic benefits study and a survey of selected PRECARN members identified the following benefits to the process control industry:

- more timely diagnosis of process faults
- more precise identification of process trends indicative of malfunctions or upsets
- more efficient, and informative methods for filtering and presenting alarms to process operators

It is expected that the implementation of APACS will result in an overall increase in the efficiency, reliability, and availability of process control facilities. The system architecture of APACS will be designed to ensure the applicability of APACS to other process plants. In particular, the feedwater system was selected for prototype development since it is representative of the complexity of process industry equipment.

The APACS project will result in a system which combines advanced AI techniques in an economically significant problem domain. Further benefits are expected through the commercialization of this technology.

FLITE

The FLITE project is a program of research in the application of artificial intelligence in the domain of aerodynamic simulation. The project is a joint effort of CAE Electronics Ltd., McGill University, and CRIM under the direction of Dr. Renato De Mori. Full-time work began in January 1986, with CRIM's participation beginning in April 1987.

The goal of the research is to build a knowledge-based system, called the FLITE System, to act as an assistant to engineers performing verification and tuning of the aerodynamic models used in flight simulators manufactured by CAE. Verification consists of comparing manoeuvres performed in the aircraft and simulator. Typically a hundred or more distinct flight-test manoeuvres are tested; the recorded data are in the form of signals. When verification encounters a discrepancy between manoeuvres, tuning is necessary. Coefficients in the aerodynamic model or simulator inputs can be modified to get acceptable simulator outputs.

Main Components of the FLITE System

The three main components of the FLITE System are the user interface, a mathematical optimization program, and qualitative reasoning. The FLITE system exploits the complementary strengths of numerical and AI techniques to solve a real-world problem.

The FLITE System searches through a problem-space of possible simulator faults and possible corrective actions. In practice, correctly determining the cause-and-effect relationships in such a highly dynamic, highly cross-coupled system is the main source of difficulty. Qualitative reasoning about causes of discrepancies is employed to guide the search by identifying "suspicious" parameters, those which can be modified to reduce discrepancies. Qualitative models based on linear differential equations are used to model time-varying signals and the relationships between them. The qualitative representation uses intervals with discrete end-points and a qualitative calculus to compute the effects of changes in parameters.

Parameter identification techniques are used to fine-tune suspicious parameters. This approach is based on the contributions of Dr. P. Belanger of McGill University, an expert in this area. The problem of tuning parameters is mapped to a mathematical optimization problem and solved by a numerical algorithm. The results of tuning are exact, specific changes to aerodynamic coefficient or inputs which minimize discrepancies. The optimization will be feasible and the results meaningful, since they are based on an approximate solution from a qualitative model.

The user interface is based on a multiwindow graphics display with a mixed-initiative dialogue style. A user can, at any time, request graphs and displays to monitor the state of problem solving. The FLITE System will present displays to the user when it considers them important and will periodically ask for user input to guide problem-solving. Users may answer immediately or postpone their response while they request and study other displays.

Research Contributions

The main contributions to AI research arising from the FLITE project have been:

- improved qualitative reasoning techniques for time-varying systems
- integration of qualitative reasoning and numerical optimization
- improved constraint satisfaction algorithms for qualitative reasoning

Two post-graduate degrees are expected as the result of work performed for the FLITE project. These are: R. Prager, CAE Electronics, Ph.D. Computer Science from McGill University, expected 1990, and J. Levine, CRIM, M.Sc. Computer Science from McGill University, expected 1990.

This research has already led to two major publications. A paper entitled "A Knowledge-Based System for Troubleshooting Real-Time Models," by R. Prager, P. Belanger, and R. De Mori, appears as a chapter in *Artificial Intelligence, Simulation and Modeling,* edited by L.E. Widman, K.A. Loparo, and N.R. Nielsen. This book was published by John Wiley and Sons in 1989. A paper entitled "Perturbation Analysis with Qualitative Models," by R. De Mori and R. Prager, was presented at the Eleventh International Joint Conference on Artificial Intelligence, held in Detroit, Michigan, in August 1989. Further publications are in preparation.

Rule-based Tools for Tactical Scenario Design

CAE designs and develops tactical mission simulation systems to meet the training requirements of a number of military customers. An important consideration for users of these simulators is the ease with which scenario designers can create and modify elements of the simulation in an off-line design environment.

The goals of the AI Group's work in tactical simulation are:

- to provide high-level tools to allow a scenario designer to translate his knowledge of tactical doctrine into a natural rule-based format
- to enhance simulation realism by modeling the decision processes of complex agents in the simulation
- to provide graphical tools for monitoring the execution of rule bases

These objectives are currently being realized as part of ITEMS (Interactive Tactical Environment Management System). In the current version an Interactive Rules Editor and Debugger can be used to create and modify rules expressing tactical doctrine for air target response and coordinated battle group maneuvers. ITEMS is being developed on IRIS workstations in C under X-windows.

Knowledge-based Techniques for Speech Recognition

CAE is investigating the use of knowledge-based techniques to increase the recognition accuracy of speaker-dependent, continuous speech recognition systems. The application area currently being explored is air traffic control, a domain characterized by a hands-busy, eyes-busy task environment.

In the operational ATC environment CAE is targeting direct voice input to operational support systems in order to enhance or replace the current keyboard/keypad entry process.

For ATC simulation CAE is exploring the combination of voice recognition and voice synthesis to help automate the simulation sessions by replacing the need for human pseudo-pilots — ATC training personnel who play the role of pilots and controllers.

Pseudo-pilot replacement will require research and technology development in speech recognition and AI. Towards this end CAE is currently investigating the development and integration of pragmatic knowledge sources to:

- dynamically constrain syntax networks using contextual and historical information
- detect errors made either by the speech recognition system, or the speaker
- correct recognition errors and provide more meaningful error messages

These efforts are being conducted as part of a contract with the Transport Canada Transportation Development Centre.

Other Activities

The AI group's activities have two major goals, to add value to existing CAE products and to improve the efficiency of CAE's operations. To help achieve these goals, specific development projects are being sought through contract work and in-house applications.

Several contract proposals submitted by CAE to customers have included AI components. These are in areas where CAE has traditional strengths, namely training systems and control systems. CAE expects to integrate AI components within its product line to meet the evolving needs of its customers.

A survey is being conducted for potential in-house projects which apply well-proven AI techniques to the solution of specific problems. Initial interest has been expressed in scheduling systems and configuration systems. Commencement of these projects is expected in early 1990.

Canadian Pacific

Canadian Pacific started using Artificial Intelligence technology in 1984. Several applications have been developed that use AI tools and methodologies. AI tools are used for prototyping, gathering and documenting expertise, and developing applications. Some applications that have been implemented at Canadian Pacific are:

- *Expert Radio Fault Diagnostic System*: This application was designed to troubleshoot 80% of the faults occurring in Motorola's Model MT500 radios. Development was started in 1984 and the rule-based system was implemented in Saskatoon in 1986.
- *Lube Oil Analyst Expert System*: This is a diagnostic application that applies the expertise of CP's lubricant specialists. Following a spectrographic (or physical) analysis of a locomotive oil sample the expert system is invoked to produce a recommendation on the status of an engine. Based on historic information from the interactive network, it recommends either normal operation, a low-level alert, or to book a unit for maintenance. The application performs the engine oil analysis for one locomotive unit in less than 30 s. It is installed at the oil control laboratories in five yards across Canada.
- *Computer Assisted Manual Block System*: This application monitors rail traffic, applies train dispatcher rules, and lists all the valid choices of where a train may be dispatched. With CMBS, dispatch personnel use an intelligent terminal to support dispatching, planning, and train control functions. CP's family of computer-assisted train dispatching products have operated with an excellent safety record since the first installation in 1985. These products control 20 000 miles of track on North American railways. Within the next year CP's products will be used to control over 30 000 miles. By reducing the work necessary to control trains CMBS helps reduce dispatcher stress and increase rail safety.
- *Home Routes*: This is a mainframe application that applies the regulations of the American ssociation of Railroads to the problem of determining when rail cars should be accepted onto CP lines and how they should be returned to the owning railroad.
- *CHAR*: This is a PC front-end where users can enter "easy to use" English queries against mainframe databases. The application currently interprets queries to a DB2 database that stores the historical locomotive maintenance data.
- *Disk Failure Recovery Advisor*: This mainframe application recommends an efficient method of restoring the contents of disk packs, to minimize service disruption. The advisory system can be used to recover from hardware failures or to relocate volumes for performance purposes.

Some applications that are currently being developed at Canadian Pacific are:

- *Locomotive Diagnostic System*: By consolidating the expertise of experienced mechanics in the CP Rail diesel shops a diagnostic application is being developed which will help troubleshoot CP Rail's 645E3 diesel engines. The expertise from four diesel shops will be consolidated to produce a standardized set of procedures. The expert system will identify all problems in an engine when it comes in for repairs, decreasing the number of repeater visits to the shop.
- *Help Desk Problem Walkthru*: When a person calls the help desk with a problem this application will lead the support person through a series of questions to be asked over the phone. The person may also be asked to perform tests and enter commands to resolve the caller's problem.
- *Work Gang Simulation*: We are developing a simulation model to aid the Engineering department in scheduling repairs on the rails and ties. The schedule must be designed to minimize train delays and reduce track maintenance expenses.
- *CHAT Enhancements*: The English language query facility will be upgraded to interface with SAS databases, IMS databases, and reports stored in our internal electronic mail system.

- *Intelligent Exception Reporting*: This application will filter out pertinent information from a multitude of reports available on our network. Individual users can specify what conditions to look for, and when these conditions are met, the data will be extracted and put on an exception report to be sent to the user.

For further information, please contact:

Candy Baginski
Project Manager
Computers and Communications
Canadian Pacific
40 University Avenure, Suite 1200
Toronto, Ontario
M5J 1T1
Telephone: (416) 863-8782

CompEngServ Ltd.

CompEngServ is a 15-year-old company dedicated to advanced computer hardware and software. Since 1983 it has been actively involved in developing AI solutions to problems.

The company's expertise in AI crosses the classical AI disciplines of expert systems, rule-based systems, planning systems, vision systems, robotics, and neural networks. The tools utilized in these domains include object-oriented programming languages, knowledge acquisition, and waterfall design techniques.

CompEngServ employees have published and/or presented over 20 papers on AI or AI development projects. CompEngServ is considered a world authority on AI and Air Traffic Control. Other domains in which CompEngServ has worked include physical oceanography, finance, diagnostics, configuration management, air quality diagnostics, and satellite data programming.

A short summary of several major projects follows.

Conflict Resolution Expert System

Contract with Transport Canada

CompEngServ is currently involved in the development of expert systems for air traffic conflict resolution. The systems (named CORES for COnflict RESolution) are being developed for the Transportation Development Centre.

These systems are based on previous work which resulted in a prototype for the resolution of conflicts involving only two planes. The Phase II system extended this work in several areas including user interface issues, consideration of multiple planes, and issues surrounding induced conflicts. The CORES system functions within airspaces where solving an initial conflict may induce further conflicts. The issues of convergence from an airspace with induced conflicts to a conflict-free airspace is central to the project. Strategies involving constraint relaxation for tactic selection were also examined.

The system was implemented using the LISP-based system builder KEE. It runs on a 386-based Unix workstation with a high-resolution monitor and mouse.

The Phase III work currently underway moves the problem into a more tactical radar environment issue of realtime performance and constraints are central to this project.

This phase is being implemented on Gensym's G2 realtime expert system tool on a networked DEC station environment.

Electromagnetic Interference Expert System

Contract with Transport Canada

The Electromagnetic Interference Expert System (EMIR) is a classical diagnostic package composed of three main functions: first a set of queries which attempt to form a set of hypotheses of possible causes, then a set of recommended tests to validate or further reduce the hypothesis set, followed by suggested remedial actions. The tests and the remedial actions all have costs associated with them and the system attempts to offer minimum cost alternatives to the operator.

A particularly interesting part of this system is the requirement for use by qualified, but not expert, technical officers in this diagnostic process.

Finally, the system must generate reports not only of the interference, but the process of solving the problems for review by the centrally located expert.

This system was implemented in Insight 2+, a Pascal-based tool. It contained 150 rules and an extensive database.

Marine Environmental Data Management System

Customer: Department of Fisheries and Oceans

The EDMAN system is an intelligent interface to archival data with a set of tools for data manipulation which are suitable for the domain of the data. This system can be further enhanced to provide many new tools for data viewing and manipulation. EDMAN services three types of users. The first is the developer who maintains the system and adds new functionality. Next is the technician who is responsible for checking the quality of data before it is entered into the appropriate database. The final user is the scientist who retrieves data from the database and performs "data snooping."

All three types of users are provided with a graphical environment which we call the workroom. The workroom contains a toolbox with a variety of "drawers" in which tools are stored. These tools can be anything from expert systems to Fourier transform routines to *x–y* plotting functions. The workroom provides a workbench where these tools can be put together to form applications. In its current implementation, data flows from one tool to the other in a sequential pipeline. However, we plan to allow the use of conditions and loops in the future, thus providing a very high-level graphical programming equivalent.

The first application built for the EDMAN environment is the Temperature, Salinity and Current Quality Control (TSAC-QC*) system. This application is an automated version of the Thermal Salinity Quality Control Manual developed by a domain expert. The manual defines a series of tests in an IF-THEN format for determining if TSAC data might contain inconsistencies or errors. Each test is implemented in TSAC-QC as a tool consisting of a set of forward-chaining production rules. The complete application consists of a compound tool produced by linking the primitive tools together. It provides expert advice to a technician who is responsible for ensuring the quality of data before it is entered into the Data Base Management System (DBMS).

The EDMAN system has been implemented in C++ on a VAXstation 2000 running VMS. CompEngServ is currently marketing this program around the world.

For further information, please contact:

David G. Bowen, P. Eng.
Director of Engineering
CompEngServ Ltd.
600-265 Carling Avenue
Ottawa, Ontario
K1S 2E1

Digital Equipment Corporation

Background on Digital and AI

Digital Equipment Corporation was an early user of expert system technology and has successfully fielded over 80 expert systems in manufacturing operations, manufacturing planning, engineering design, financial analysis, order and product configuration, distribution, field service, sales operations, and marketing data analysis.

Digital maintains a vital and aggressive program of expert system development. Currently, Digital realizes annual savings of $200 million from deployed expert systems, and this figure is growing at 40% per year.

Digital has a staff of over 600 AI specialists worldwide, integrated into the company's various manufacturing, engineering, information systems, product development, and professional services organizations. Specialists in expert system development reside in most major countries in which Digital does business, including Canada, the United States, Japan, Europe, and Australia. The development and deployment of AI technology at Digital is an international effort.

The hub of AI activity is Digital's AI Technology Center (AITC) in Marlborough, Massachusetts, which houses specialists whose mandate is to provide worldwide support for AI activities and AI professionals in the field. There are groups in the AITC which are focused on the following:

- *Product development*: development and maintenance of commercial expert system tools such as OPSS, VAX LISP, and VAX Decision Expert, as well as internal tools tailored for specific application areas, such as diagnostics and configuration.
- *Internal applications*: development of expert system applications (such as XCON, XSEL, MOC, CANBUILD, and others) to support the operations of various Digital business units.
- *External consulting*: provision of expert system consulting services to Digital's major accounts. Specialists in this group have expertise in specific industries, such as process manufacturing, financial services, and telecommunications, and assist local account teams in delivering such services as application analysis, prototyping, system development, management consulting and technology transfer.
- *AI research*: research in leading-edge applications of artificial intelligence and liaison with major universities and research consortia.

Development of AI Specialists

Early on, Digital recognized the need for a large number of skilled knowledge engineers and realized that demand would outweigh supply. Digital addressed this problem by developing a core curriculum for training knowledge engineers and enroling suitable applicants from Digital business units with expert system applications to develop. In this manner, Digital was able to bootstrap its existing investment in highly trained software professionals and disseminate AI expertise throughout the company.

Digital's knowledge engineering training program began in 1981; it consists of formal and apprenticeship phases. The formal phase consists of lecture and laboratory work directed by Digital staff and instructors from cooperating universities, such as MIT, Carnegie-Mellon, and Harvard. University instructors include Randall Davis and Patrick Winston, MIT, Marianne LaFrance, Boston College, and Dorothy Lennart-Barton, Harvard School of Business. During the apprenticeship phase, the student develops a prototype expert system application under the supervision of an experienced knowledge engineer. He returns to his business unit with a sound foundation in AI technology and a prototyped application with an implementation plan to put it into production. Over 450 students have graduated from Digital's Knowledge Engineering Curriculum.

In 1987, Digital's program for developing knowledge engineers was made available to selected customers under the title "The Digital Business Fellowship in Artificial Intelligence" (AI Fellowship Program). Employees of more than 80 Digital major accounts have graduated from the 6-month AI Fellowship Program, including 14 from Canada.

Focus on Large-scale Integrated Applications

Digital's focus is on large-scale, highly integrated expert systems deployed in a distributed networked environment. In a typical Digital expert system project, 20–30% of the project team will consist of AI specialists. The remainder will bring expertise in other technologies, such as database, networking, imaging, and industrial automation to the problem.

As an example of this approach, Digital's XCON/XSEL expert systems, which configure computer systems according to customer requirements, are implemented across Digital's worldwide network of 80 000 users. These systems, critical to the operations of the company, are used by every sales representative worldwide as the sole means by which a quotation can be generated and an order submitted. The deployment of these systems across a distributed network enables centralized administration of configuration, product, and pricing information, and ensures that changes will be reflected immediately at all Digital offices throughout the world.

Cooperative Projects in Canada

Digital has engaged in a number of cooperative expert system projects with Canadian customers. These include:

- set-up assistant for a cold reduction mill
- alarm processing system in a large power network
- expert system for predicting severe convective storms
- scheduling system for hot metal distribution system
- computer-aided design assistant for injection molds
- financial underwriting expert system
- inventory management/order scheduling expert system
- on-site AI Fellowship program
- petroleum products pipeline scheduling expert system

Future Directions at Digital

Recently, Digital has been working toward the goal of reducing the cost of expert system applications by reducing the development effort involved. Digital recognizes that the market for "off-the-shelf" expert system applications will grow dramatically, but that a high degree of customization will be required to transform prepackaged applications into business solutions.

One direction that Digital is taking is to identify solutions that Digital has developed internally or for customers where the problem addressed is common to all or most companies in that industry. Typically, the core of the solution is applicable to any company in the industry, and a solution for a particular company can be achieved by surrounding the core technology with specific consulting services. At this point, Digital has developed programs for supplying core technology enshrouded in consulting services for the following applications:

- cross-selling in a retail banking environment
- oil and gas pipeline scheduling
- property and casualty underwriting
- marketing data analysis

A second direction for Digital arose from the recognition that expert system solutions are often best classified according to the nature of the problem (diagnostics, configuration, scheduling, etc.) rather than by industry. Digital has developed internal tools and techniques for solving particular kinds of problems, and these are being packaged with consulting and training to enhance a company's productivity in developing a custom solution. Currently, Digital has packaged tools and services for dealing with two problem areas: diagnostics and configuration.

Finally, Digital is exploring emerging areas of commercial artificial intelligence by engaging in cooperative projects with customers. As an example, Digital's Machine Learning Group is working with selected customers to apply neural network, adaptive pattern recognition, and parallel associative memory technology to problems such as image recognition, text retrieval, and process and quality control.

A Decade of AI Leadership

For the past decade, Digital has assumed a leadership role in the application of artificial intelligence technology to critical business problems. Digital is well positioned to continue this leadership throughout the next decade, in Canada and internationally.

For further information, please contact:

Gordon Graham
Marketing Manager, Artificial Intelligence
Digital Equipment of Canada Ltd.
4110 Yonge Street
Willowdale, Ontario
M2P 2C7
Telephone: (416) 730-7308

Hydro-Québec

Large electric utilities can benefit significantly from expert systems technology.

If power systems are to achieve a high degree of reliability and efficiency, they must use the most powerful and promising technologies available. In this regard, artificial intelligence, more specifically expert systems, could have interesting potential. The technology of expert systems is worth investigating because most power system activities are knowledge-intensive.

According to Harmon and King, the authors of the well-known book *Artificial Intelligence in Business*, two waves of expert system applications will take place. The first will last until the beginning of the 1990s and will be dominated by medium-sized applications designed to solve specific problems. Proliferation of this type of expert system will have the merit of fostering a favorable attitude toward this technology and contributing to the training of large numbers of knowledge engineers. It will also reap the first economic rewards for this new technology.

The second wave, which will feature much larger systems, will start to have an impact early in 1992 and will peak at the turn of the century. These expert systems will be far more apt to reproduce human reasoning, to the point that they will even rival human experts at times. This second wave will herald a veritable revolution likely to shake the very foundations of present-day methods of processing and disseminating information, as well as the decision-making process.

This scenario has been confirmed by market development projections. A prestigious consulting firm from the United Kingdom predicts a 6.6-fold increase in the market between 1986 and 1992 for Canada and the United States alone, or a sales volume of US $1.9 billion in 1992 for expert-system software and hardware.

Researchers at Carnegie-Mellon University and EPRI, a group of Japanese utilities, and Hydro-Québec researchers have attempted to assess the potential expert systems offer electrical utilities. The following conclusions can be drawn from their findings:

- Since the potential applications of expert systems in the electrical industry are determined by its know-how, the potential benefits are virtually unlimited.
- Current applications lack coherence and do not follow any coordinated plan.
- Most applications are experimental.
- The areas of application explored so far represent only the tip of the iceberg.

The existing applications for expert systems at electrical utilities are found in the following areas: planning, system diagnosis, equipment diagnosis, engineering and design, decision-making, and training.

Hydro-Québec System

In January 1988, Hydro-Québec's installed capacity was 25 000 MW, of which 93% was hydro power. Its total sales in 1987 were $5 billion. The transmission system consists of 33 000 km of lines, of which 33% is 735-kV lines. The distribution system consists of 99 000 km of lines. Projects now under way will add 3000 MW.

The Project Selection Process

Hydro-Québec's process of selecting candidate expert system projects begins with an analysis of the problems involved in light of established criteria. Two basic questions must be asked:

1. Is the expert system possible?

2. Will the expert system be useful?

To arrive at logical answers, two evaluation grids have been developed, the first designed to eliminate unsuitable candidates, the second to rank those that remain according to a scale of relative values.

First Grid: Is the Expert System Possible?

- Is the problem essentially heuristic and difficult to solve by conventional algorithmic methods?
- Is local expertise available?
- Is that expertise acknowledged by the users?
- Is knowledge engineering available?
- Does a data processing infrastructure exist?
- Is the size of the system reasonable?
- Is the degree of difficulty manageable?

Second grid: Will the Expert System be Useful?

- Is the expert system justified? Does it offer significant, measurable, economic advantages? Is it in keeping with the utility's development plan?
- Is the expert system necessary? Is there a need to transfer the expertise to a larger group of users? Is expertise disappearing? Is there a need for training? Has a need been identified for expertise in a hazardous environment?

Hydro-Québec's Expert Systems Projects

TRANSEPT: Expert System for the Preliminary Design of a Transmission System

Planning of transmission system facilities is based on an immense body of knowledge, which is scattered throughout the utility. TRANSEPT is a preliminary design tool that can establish a score of alternatives. It takes into account the number of lines, voltage, compensation, and system stability margin, while at the same time allowing the user to assess the financial impact of these various parameters.

The purpose of TRANSEPT is to select a reduced number of alternatives that help the designer with the various design stages. The block diagram of TRANSEPT is given in Figure 4. This expert system can integrate algorithmic models and simulation models with the expert's knowledge and experience.

The explanations provided by TRANSEPT are extremely useful in securing user acceptance of proposed alternatives, as well as training new experts in system design. The TRANSEPT prototype was completed in 1988 and is now being validated.

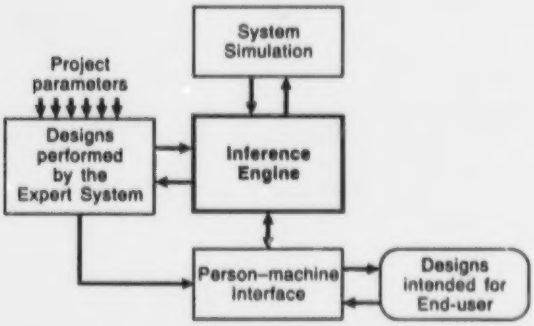

Figure 4. Block diagram of the TRANSEPT expert system.

SEPIDE: Expert System for Interpreting Events

Hydro-Québec has set up a large database designed to detect and record in detail all changes in status occurring on the transmission system. After an event, the changes in status recorded in this database are analyzed by the SEPIDE expert system which produces an explanation of the event which makes it possible to evaluate the situation. This activity is essential if the utility is to maintain the high level of reliability expected by its customers. Moreover, the maintenance schedule depends entirely on the results of this type of analysis.

The purpose of the SEPIDE project is to automate realtime or off-line interpretation of all transmission system events forming part of normal system operation (70% of them being automatic control operations). Figure 5 shows the database that comprises the chronological records of events (CRE) and constitutes the main input of the system. The SEPIDE prototype was developed in 1987 to establish the principles underlying this system and to demonstrate its feasibility. A SEPIDE pilot project has been installed and is now being validated.

Expert System for Alarm Processing

The task of subtransmission system operators is becoming increasingly demanding as a result of the greater complexity of facilities as well as a substantially greater volume of information supplied by each installation. When equipment failures occur, these operators must quickly analyze the alarms and changes in status to establish fault

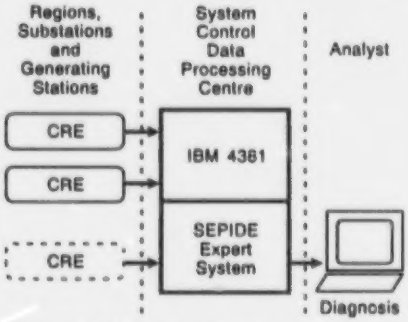
Regions, Substations and Generating Stations
System Control Data Processing Centre
Analyst
CRE
CRE
CRE
IBM 4381
SEPIDE Expert System
Diagnosis

Figure 5. Block diagram of the operating environment of the SEPIDE expert system.

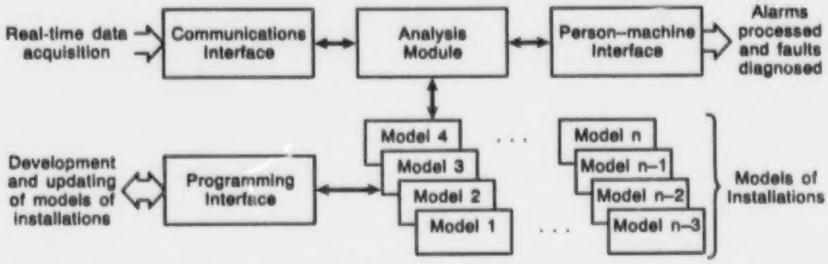

Figure 6. Block diagram of the alarm processing expert system.

diagnoses. These diagnoses enable them to take the necessary measures for power system restoration.

Operators are, however, frequently faced with complex analytical problems that require several minutes of reflection. These delays affect the quality of service because they prolong the outages. Moreover, the outages may increase in number if appropriate action is not taken swiftly.

This alarm processing expert system shown in Figure 6 is designed to improve the processing of alarms and changes in status received by subtransmission system operators. On the basis of the alarm sequences and changes in status, the expert system automatically establishes concise, realtime diagnoses identifying the origin and consequences of a disturbance. These diagnoses enable subtransmission system operators to accelerate their analysis and more effectively manage the information that reaches them. The pilot project was completed in March 1989. It is now processing alarms from a dozen facilities in the Maisonneuve Region in order to ensure its validation in a real operating environment.

SEDA: Expert System for Equipment Diagnosis

The SEDA project reflects Hydro-Québec's efforts to discover better ways of determining or diagnosing the in-service state of its equipment. For this purpose, it has implemented preventive maintenance, management, and monitoring systems in recent years. The advantages of these tools appear to be obvious. Today, the decision to develop an expert system for diagnosis of large power transformers is based on far-reaching experience in solving specific problems related to this type of equipment using the rules deduced and facts observed over long years of operation. This expertise must now be directed to the level of foremen and crew leaders.

The SEDA expert system, shown in the block diagram in Figure 7, is designed specifically for diagnosing phenomena that occur in large power transformers during automatic tripping.

SUBAREX: Expert System for Dam Surveillance

Dam surveillance is the responsibility of regional personnel. Design and behavior-analysis expertise is, however, spread among the various units of the utility, including those at Head Office. The result is on-site inspections followed by multiple exchanges of information, and then evaluations, consultations, and proposed solutions.

SUBAREX (see Figure 8) is a tool for diagnosing the safety of earth dams on the basis of facts observed and sounding data. This tool, which initially will be intended for inspectors, will be used by regional personnel to make decisions as swiftly as possible within a high-expertise technical framework. A subsequent version will be intended for engineers to facilitate their evaluations of the structural and operational safety of dams and to monitor changes in their behavior. A preprototype used to demonstrate the feasibility of the system is now being developed.

ENERGEX: Expert System for the ENERGIE III Model

The ENERGIE III model is used to simulate the operation of hydroelectric complexes.

Before the model can be used, however, it must be calibrated, a task that demands a high level of expertise and effort.

The purpose of the ENERGEX expert system, shown in the block diagram in Figure 9, is to concentrate varied and intuitive experience in the form of precise rules that facilitate the use of the ENERGIE III model, while at the same time ensuring a quasi-optimum solution in terms of firm

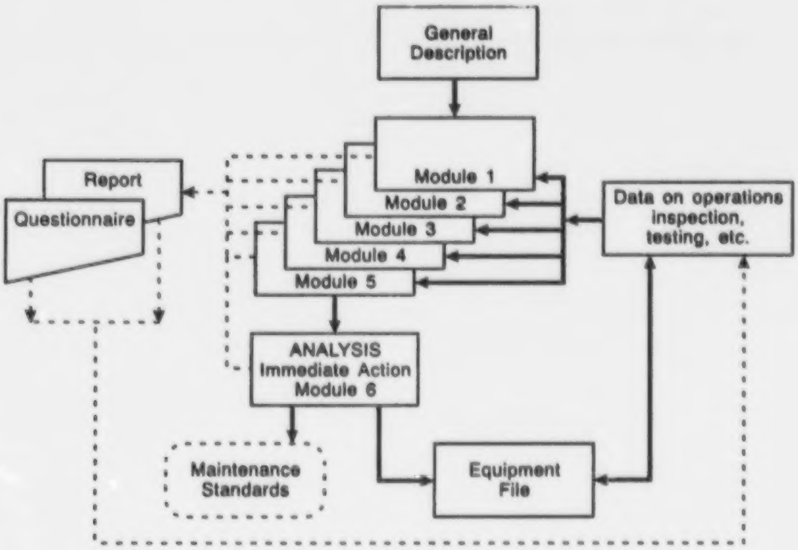
General Description
Report
Questionnaire
Module 1
Module 2
Module 3
Module 4
Module 5
Data on operations inspection, testing, etc.
ANALYSIS Immediate Action Module 6
Maintenance Standards
Equipment File

Module 1: Equipment Status
Module 2: Inspection Data
Module 3: Tests—Insulating Fluids
Module 4: Tests—Equipment
Module 5: Restoration

Figure 7. Block diagram of the SEDA expert system.

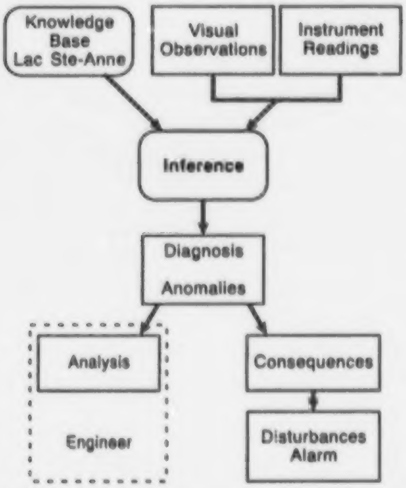
Knowledge Base Lac Ste-Anne
Visual Observations
Instrument Readings
Inference
Diagnosis
Anomalies
Analysis
Engineer
Consequences
Disturbances Alarm

Figure 8. Block diagram of the SUBAREX expert system.

energy generated. ENERGEX will bring about substantial savings by accelerating the design cycle of hydroelectric complexes. A preprototype has been evaluated and validated, and the prototype is now being developed.

EXSEMBLE: Expert System for Producing Substation Layouts

Substation design involves extremely diversified knowledge, which has been partly standardized but is found throughout the utility. In the first stage, EXSEMBLE unifies this expertise and uses computer-aided design software (AutoCAD) to produce a type-II 120-25-kV substation layout.

EXSEMBLE, which is shown in the block diagram in Figure 10, constitutes an interactive link between the knowledge base and the computer-aided design software.

As the first prototype of this kind, the EXSEMBLE project opens a very promising avenue for intelligent support of computer-aided design in all engineering activities using a high concentration of standards.

SAGE: Expert System for Management of Energy Exchanges

Management of energy exchanges with neighboring power systems involves a large quantity of information that requires realtime processing by

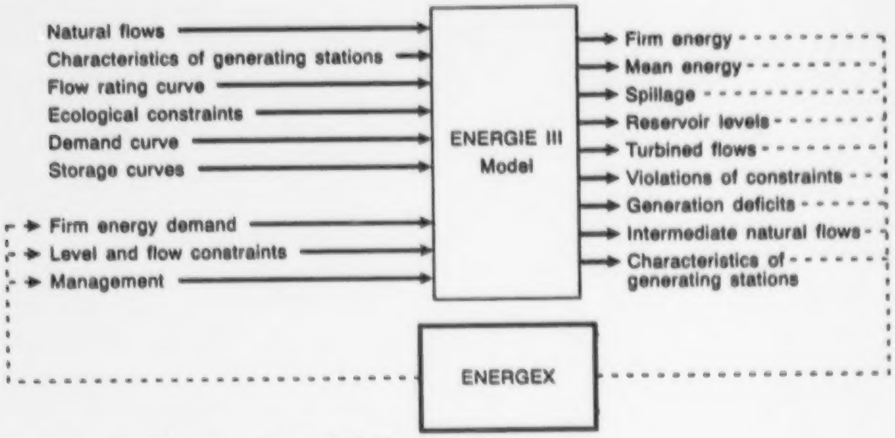
Natural flows
Characteristics of generating stations
Flow rating curve
Ecological constraints
Demand curve
Storage curves
Firm energy demand
Level and flow constraints
Management
ENERGIE III
Model
Firm energy
Mean energy
Spillage
Reservoir levels
Turbined flows
Violations of constraints
Generation deficits
Intermediate natural flows
Characteristics of generating stations
ENERGEX

Figure 9. Block diagram of the ENERGEX expert system.

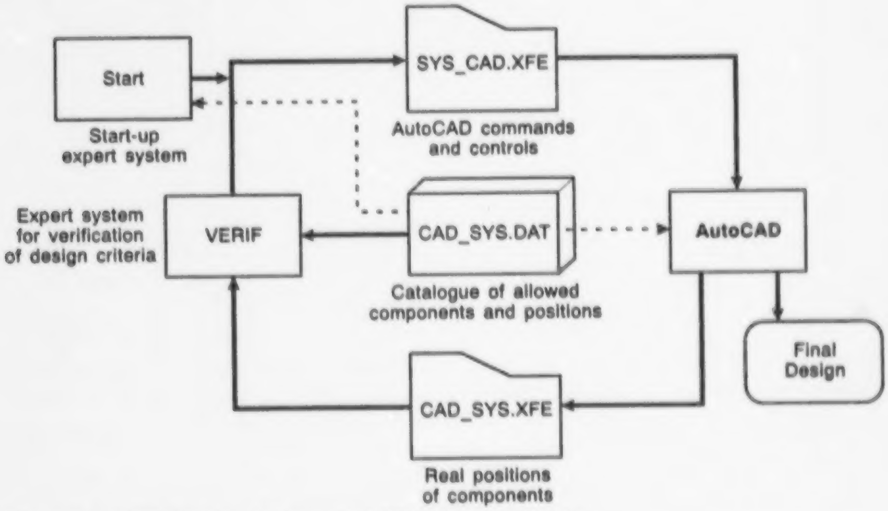
Start
Start-up expert system
SYS_CAD.XFE
AutoCAD commands and controls
Expert system for verification of design criteria
VERIF
CAD_SYS.DAT
Catalogue of allowed components and positions
AutoCAD
Final Design
CAD_SYS.XFE
Real positions of components

Figure 10. Block diagram for integrated EXSEMBLE application.

dispatchers. The SAGE expert system will support energy-exchange management while simultaneously maximizing sales revenues and minimizing purchasing costs and contractual penalties. Moreover, SAGE will standardize energy-exchange approaches and strategies, while ensuring a high level of flexibility with respect to new constraints, strategies, and calculation tools.

The SAGE expert system, shown in the diagram in Figure 11, will use realtime information from the System Control Centre (SCC) and a production-management plan established by the NUMA mathematical model, as well as rules based on energy-exchange contracts and Hydro-Québec management standards. SAGE will establish the program for final exchanges on the basis of this information and these rules.

Development Tools

One of the main problems to be faced in building an expert system is to select the appropriate tools, since their features have a considerable influence on the quality of the system developed, especially its user-friendliness and performance.

To allow the designer to make an enlightened judgment, Hydro-Québec has developed a

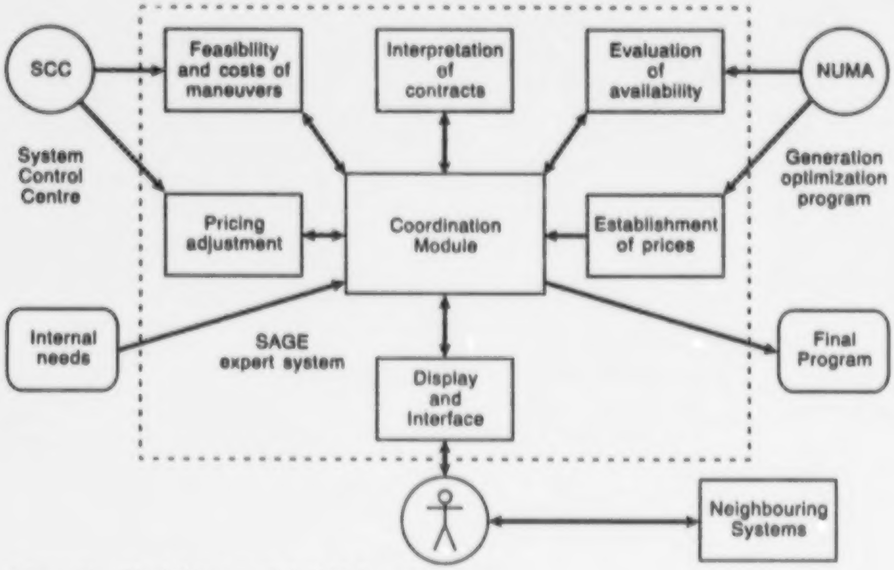
SCC
System Control Centre
Feasibility and costs of maneuvers
Interpretation of contracts
Evaluation of availability
NUMA
Generation optimization program
Pricing adjustment
Coordination Module
Establishment of prices
Internal needs
SAGE expert system
Display and Interface
Final Program
Neighbouring Systems

Figure 11. Block diagram of the SAGE expert system

two-tier selection grid, one concerning the knowledge engineer and the other the end-user.

The first grid, which concerns the knowledge engineer, is used to compare the properties of the problem to be solved with the characteristics of the proposed tools. It affects the following elements:

- knowledge representation: production rules, predicate logic, semantic networks, frames
- control: forward, backward, and mixed chaining, bidirectional chaining, blackboard architecture
- logic: use of variables, negation processing, etc.

The second grid, which concerns the user of the expert system, focuses on the following:

- user-friendliness
- explanations: explanation of reasoning and justification of conclusions
- interfaces: graphic, databases, signals, etc.

The characteristics of the tools considered are evaluated by means of performance coefficients established using these grids. These coefficients are then used to compare the candidate tools with the characteristics of the problem to be solved. At the same time, Hydro-Québec endeavors to be economical with respect to the number of tools it employs in order to maintain reasonable budgets and limit training requirements for knowledge engineers. The price/performance graphs published by Harmon, Maus, and Morrissey in their book, *Expert Systems — Tools and Applications*, albeit very preliminary, are a good resource for tool selection.

Table 7 shows the tools selected for the projects referred to in this paper. The reader will notice a polarization around certain tools. Subsequent experience, project evolution, user reaction, and the advent of new tools on the market could modify this pattern. As well, developments in problem knowledge between the feasibility study and implementation of the finished product may entail a modification of the tool selected, in particular after the prototype stage. Moreover, in certain cases, different environments may be required for development and for operations.

Table 7. Tools used for the expert systems developed by Hydro-Québec.

Tool	Expert system
Goldwords	ENERGEX SUBAREX TRANSEPT
OPS 83	Alarm processing system
Rulemaster	EXSEMBLE SEDA SEPIDE

All these expert systems were developed on 386-type personal computers, with the exception of the Alarm Processing Expert System, which was developed on MicroVAX-GPX, as well as SEPIDE, which was developed on HP 9000.

Organization

In his book, *Expert Systems in the Workplace*, O'Farrell cites four criteria for ensuring success in the development and implementation of expert systems in a utility environment. These criteria are:

1. perceptive selection of projects and project specifications with well-defined objectives
2. emphasis on domain experts
3. presence of a "champion" of the cause
4. appropriately structured organization

Hydro-Québec satisfied criteria (1) and (2) from the outset by developing the above-mentioned selection grids and by creating multidisciplinary teams, each comprising the domain expert, the end-user, and the knowledge engineer, who report to the same authority. Moreover, special attention is given to developing clear and well-documented mandates involving precise objectives.

A Task Force on Expert Systems has been set up. It comprises highly motivated individuals from various administrative units of the utility. This horizontal structure facilitates communications and controls interdepartmental and professional competition. Thus criteria (3) and (4) are also satisfied.

Hydro-Québec's Task Force on Expert Systems, the structure of which is shown in Figure 12, is responsible for initiating and taking concerted action to introduce expert-system technology to the utility's main areas of activity. Its mandate includes the selection, development, implementation, and marketing of expert systems in areas related to generation, transmission, distribution, and use of electricity.

The Task Force reports to a Steering Committee comprising senior management representatives from the departments involved in activities related to expert systems, under the responsibility of the Vice-President, Development and Technology Marketing.

A Technical Committee, consisting of managerial personnel from the main units involved in expert-system projects contributes to the orientation of the Task Force's general activities, while current activities fall under the responsibility of a Project Manager.

This flexible structure is designed to establish horizontal liaison allowing dialogue between the various participants under the umbrella of the hierarchy without necessarily passing through it.

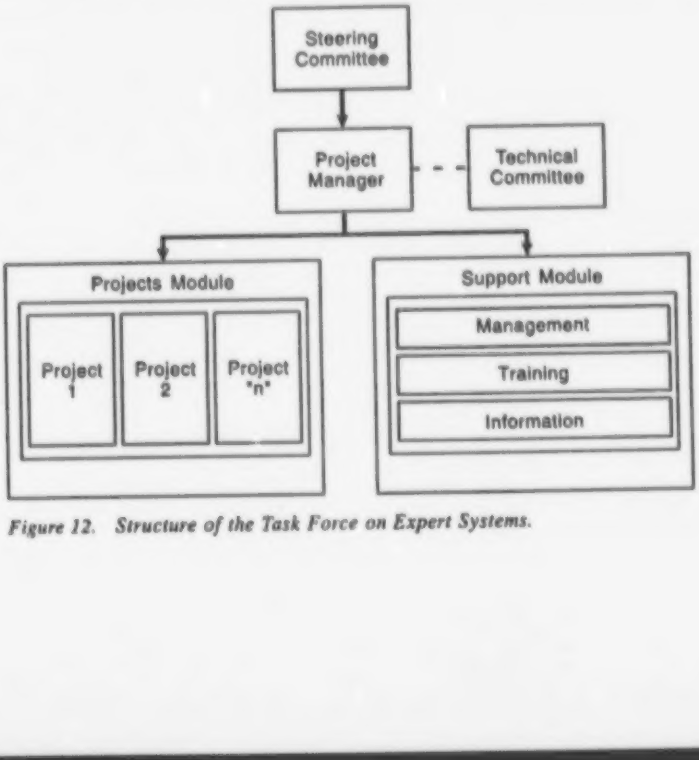

Figure 12. Structure of the Task Force on Expert Systems.

These aspects are reflected in the organization of every compartment of the Task Force. Its basic structure, for example, is represented by the Projects Module and the Support Module, in which all the human resources allocated to a project or a support activity are concentrated, regardless of the department to which they belong.

The Projects Module comprises all activities related to expert-system projects. These are carried out by teams under the management of a coordinator. The teams comprise knowledge engineering specialists, domain experts, and computer specialists, who are assigned to the projects at the time the terms of reference are negotiated for the studies to be conducted.

The Support Module covers training, project management (planning, follow-up, cost control, etc.), and dissemination of information on expert-system technology throughout the utility. The participants in these activities are delegated by their respective units and have specific mandates.

Conclusion

In conclusion, it is clear that large electrical utilities can only benefit significantly from expert-systems technology. In Hydro-Québec's specific case, this conclusion is founded on the following factors:

- There exists a large number of potential expert-systems applications in the electrical utility environment.
- Substantial expertise usually exists in areas related to electric power systems.
- Significant scientific data processing resources exist which can potentially be retrained to become knowledge engineers.
- The need to transmit expertise to the coming generation is increasingly pressing.
- This technology, over the next decade, will affect most activities related to information processing.

Acknowledgments

Thanks are due to Toby Gilsig, Elie Saheb, Roger Beauchemin, Paul Chicoine, Oscar Dascal, Sergio del Pozo, Réal Galipeau, Jean-Louis Jasmin, Miguel A. Marin, Donald McGillis, Jean-Marc Pelletier, and Gia Tong Vuong for their invaluable assistance with the preparation of this article.

For further information, please contact:

Radu Manoliu and André Lemire
Hydro-Québec
455 René Lévesque Boulevard West
Montreal, Quebec
H2Z 1Z2
Telephone: (514) 289-3877
Fax: (514) 874-0810

Artificial Intelligence Work at MacDonald Dettwiler and Associates Ltd.[1]

By Max Krause[2]

Introduction

MacDonald Dettwiler and Associates Ltd. (MacDonald Dettwiler) is a world leader in the technologies of remote sensing, particularly in the reception and processing of imagery from space. MacDonald Dettwiler has developed and installed systems worldwide for the reception and processing of data from earth resource satellites such as LANDSAT and SPOT, from meteorological satellites in both geosynchronous and near-polar orbits, and from synthetic aperture radar (SAR) satellites. Since its founding in 1967, MacDonald Dettwiler has striven to anticipate, to develop, and to use state-of-the-art technology in its products. AI is one of those technologies which appears to be poised to become an essential part of any leading-edge company's tool kit, as it has done at MacDonald Dettwiler.

Applications of AI have become an increasingly important part of MacDonald Dettwiler's research program over the past few years. With a current complement of five AI workers, the knowledge-based systems area comprises a significant and growing fraction of our research staff and research effort.

Two of the projects undertaken to date, PREFORMA and SWIFT, have been in one of the company's traditional application areas, meteorology. A third project, FEX, is applied to the problem of mapping from space, which is becoming an increasingly important area of business for MacDonald Dettwiler. The latest project, STAR, applies MacDonald Dettwiler's expertise to a new domain, maintenance of electro-mechanical systems. Although the current application domain of STAR is terrestrial, we anticipate the need for this type of expertise in space as well as in industries on earth. As a member of Canada's Space Station Project team MacDonald Dettwiler is particularly concerned with fostering and enhancing Canadian skills in AI for potential future application in that project.

The following short accounts summarize some of the AI work undertaken in the recent past at MacDonald Dettwiler.

PREFORMA

Participants: Guy Druce, Max Krause,
Peter George

The PREFORMA (PREcipitation FORecasting Meteorologist's Assistant) project was undertaken to test the feasibility of achieving expert level performance by a knowledge-based system in the field of site-specific short-term weather prediction. The approach taken was to model the conceptual categories and the inference methods applied to this problem by human weather forecasters. This approach was motivated by our perception that this problem was being solved with reasonable success by human forecasters in day-to-day operations, and that a numerical approach to the problem demanded resources that were beyond the means of most potential users of such a system. In order to limit the complexity of the problem, we took the advice of the Atmospheric Environment Service and restricted it to 12-hour prediction of the occurrence, phase, timing, and intensity of precipitation for a single site.

After about two person-years of effort, PREFORMA has a very good user interface which supports the analysis of weather maps in a way that would speed up the analytical work that an operational forecaster performs. The meteorological reasoning component of the system can perform the automatic projection of future positions of cold and warm fronts using rules of thumb employed by forecasters, although with lower accuracy than human forecasters. This discrepancy in performance is due to the difficulty of performing on the computer the complex visual pattern recognition tasks that human forecasters use to project the future positions of weather features.

SWIFT

Participants: Max Krause, Peter George (MacDonald Dettwiler), Doug Konkin, Ken Gamble, Julia Driver (Alberta Research Council)

The SWIFT (Severe Weather Intelligent Forecast Terminal) project is a joint effort by MacDonald Dettwiler and the Alberta Research Council for the Atmospheric Environment Service of Environment Canada. SWIFT is a prototype system designed to provide a forecaster with a workstation which aids in the tasks of organizing and interpreting the large amounts of information relevant to making same-day predictions of severe weather in the Canadian prairie provinces. The system makes use of both a statistical index of potential convective activity (the Sc4 Index developed by Geoff Strong of the Alberta Research Council) and knowledge of observed weather features to produce predictions of the location and direction of motion of areas where severe convective weather is likely to occur within the current day.

---

[1] Reprinted with permission from *Canadian Artificial Intelligence*, July: 23–27; 1988.

[2] Max Krause is an AI researcher employed by MacDonald Dettwiler and Associates. He holds a B.Sc. in Psychology and an M.Sc. in Computer Science.

The "intelligence" of SWIFT lies in several areas:

- a module which parses the remarks section of SA reports to extract information about weather events such as cloud and precipitation
- a module which uses knowledge of observed weather features and events to refine the interpretation of a statistical predictor of storm initiation in the eastern foothills of the Rocky Mountains (the Sc4 index)
- a subsystem which captures an internal (frame-based) model of a weather chart analysis of features and fields — the results of subjective or objective analyses are thus available to other subsystems which can use such knowledge to reason about future developments
- spatial (geographic) and temporal reasoning ability

SWIFT is implemented in InterLISP-D, LOOPS, and OPS4 on a Xerox 1186 LISP computer. The reception and formatting of incoming data and the calculation of the Sc4 index is done on a host computer (currently a DEC VAX); the workstation software accesses the host disk for data as needed.

FEX

Participants: John Zelek, Ben Yee

FEX (Feature EXtraction) is a knowledge-based system for automatic extraction of planimetric features from remote sensing imagery. Topographical planimetric features include natural surfaces (rivers, lakes) and manmade surfaces (roads, railways, bridges). In conventional planimetric feature extraction, a photo interpreter manually interprets and extracts features from imagery on a stereoplotter. Visual planimetric feature extraction is a very labour intensive operation. The advantages of automating feature extraction include savings in time and labor, accuracy improvements, and planimetric data consistency.

Present research efforts at MacDonald Dettwiler have concentrated on the extraction of road-like features in SPOT imagery. Numeric techniques by themselves would not achieve the consistency and accuracy of human photo interpretation. In order to detect and locate road features, FEX incorporates heuristic knowledge of photo interpretation. Some of the knowledge associated with roads is as follows:

- Spectral: Road segments in the imagery have higher spectral contrast than the immediate surrounding neighbourhood. Road segments tend to have a local uniform spectral distribution.
- Geometric: Parallel line structures give hints on the location of potential road segments. Locating parallel lines from the detected edges in the imagery requires the extraction of a hierarchical network of abstract geometric structures.
- Structural: Roads form networks and intersect other road segments.
- Contextual: Roads are not located in water. Houses and buildings are usually located near roads. A bridge is a road segment that crosses over a water body.

FEX combines techniques from image processing, remote sensing, and artificial intelligence for automatic feature extraction. The FEX prototype has incorporated some of the spectral and geometric knowledge. This prototype has been able to reduce commission errors and omission errors to less than 30%. Research indicates that the error levels can be reduced even further by including structural and contextual knowledge.

Numerical processing in FEX is handled by VAX Pascal, while symbolic data manipulation is managed by Quintus Prolog. FEX is implemented on a DEC microVAX.

STAR

Participants: Timothy Bult, Max Krause, Randy Roesler, Daphne Hassner, Peter George (MacDonald Dettwiler), Richard Rosenberg (University of British Columbia), Frits Swinkels (Simon Fraser University), Eric Manning (University of Victoria), Brent Sauder (MacMillan Bloedel Ltd.), Dennis Schwab (Finning Caterpillar)

The STAR (STrategic Advanced Research) project is a cooperative university/industry effort, formed under the leadership of the British Columbia Advanced Systems Institute, to carry out a knowledge-based system (KBS) project in conjunction with MacMillan Bloedel, British Columbia's largest forest industry corporation. The project aims to build a KBS for detecting and predicting faults in large off-highway logging trucks and planning allocation of resources to improve the availability of this equipment. A major catalyst for the development of this type of technology is the Canadian Space Station Program (CSSP), which is anticipated to require extensive expertise in AI technology for development of the Mobile Servicing System (MSS).

MacMillan Bloedel has installed devices on some of its trucks to record 25 types of data at regular intervals while the equipment is in operation. The STAR project will develop a KBS to analyze this collected data, detect performance degradation, predict impending fault states, and diagnose the cause of detected faults. The KBS will make qualitative judgements on the seriousness of actual or anticipated problems. The project is also investigating the possibility of planning preventive maintenance work to maximize the availability of the equipment.

In order to perform this function, the KBS will be given observations collected by either a human user or the automated monitoring system. It will then monitor these observations for symptoms, form hypotheses about the probable causes of symptoms, predict their near- and long-term effects, and recommend possible courses of action for alleviating them. The system's reasoning will be based on an internal model of the structure and function of the device and knowledge of relevant diagnostic, repair and planning strategies.

New Brunswick Power Commission

Introduction

The objective of this summary is to give a brief overview of expert systems applications within New Brunswick Power. All of these projects are directed toward the solution of specific problems encountered by the utility. In all of the projects, the utility direction and overview was in place. Some of the work, however, was carried out by agencies external to the utility such as the local university.

Diagnosing Conditions in Governors

This project was originally structured jointly by the System Performance Group and the university. It continued in this mode with overview and the utility perspective provided by New Brunswick Power and the theoretical work carried out at the university.

Objective

The objective of this work was the design and implementation of a realtime knowledge-based controller which would minimize the effect of failure-induced response in electrohydraulic governors for a hydro generating unit. Specifically the failures considered were as follows.

- Controller parameter settings did not agree with the generator loading conditions.
- Generator speed sensor output was stuck at zero.
- Failure occurring in the actuator driving the wicket gates.
- Mistuning of the integral gain of the PI governor.

The environment used was Prolog.

Unit Commitment Alternatives

This project evolved from consideration of fairly extensive modifications to an operational hydro-thermal scheduling and river modeling software package. It was decided to investigate the alternatives for the front end unit commitment module.

Objective

The objective of this work was to examine rule-based alternatives in unit commitment. One of the widely used methods of stacking units is called the priority ordering method, and one of the bases used for ordering is average unit production cost. As part of the modifications to the software package, it was considered desirable to examine alternatives for the existing front end unit commitment model.

Approach

The thermal input-output characteristics were taken in the usual quadratic term. Based on experience a look was taken at the various factors which went into the cost of production for a unit. Ultimately some 20 rules were structured, first based on average production cost and then on no load cost.

The language used in this application is the if-then-else of Fortran.

Electrical Power Distribution Applications

This project was directed at an evaluation of the feasibility of using knowledge-based system technology in the power distribution operations area.

Objective

The objective of this work was to examine the technical feasibility of applying expert systems in the automation of some operational aspects of the distribution system. As part of the objectives, a list of the technical requirements was generated. A further stipulation was that the system used should be easily interfaced with the corporate database which is on the IBM mainframe.

Approach

Fault diagnosis and power restoration were determined to be major areas of concern. The first development system considered was the Expert System Environment which runs on the IBM mainframe. The second development tool considered was GoldWorks. It is a LISP-based system which runs on an IBM-PC AT or compatible computer. The third environment considered was the Automated Reasoning Tool. It runs on the Symbolics 3600 series LISP minicomputers, and also on the Sun workstation. The final environment evaluated was Knowledge Tool which is an IBM mainframe system.

The results of this work pointed toward the Automated Reasoning Tool and Knowledge Tool as systems which met the technical requirements for implementation on a full-scale distribution network.

Test Data Development

This project is directed at development of an expert system for automatically generating test data for control programs.

Objective

The objective of this work is to develop a prototype expert system which automatically generates test data for a testing program. Specifically the objective is to automate the generation of test data used to validate the response of the Moderator Temperature Control Program.

Approach

Control programs are the computer software which drive the digital controller in a nuclear power plant. Control programs are currently tested using judiciously selected sets of data. The prototype expert system will, for a selected control program, duplicate the set of test data that is currently manually prepared. The prototype expert system will capture the rules and data that define the decisions made by the Moderator Temperature Control Program. The system will also incorporate additional rules to capture the expertise or practical skills that an engineer applies to software testing. An expert system shell called CLIPS will be used to develop this system.

Automaric Selection of Channels for Refuelling

This project was directed at providing assistance to the fuelling engineer in a Candu plant in the selection of channels for refuelling.

Objective

The objective of this work was to develop a knowledge-based system which would assist the fuelling engineer in the selection of channels for refuelling within the next few days of operation. At present this is based on experience and human judgement, and it involves close scrutiny of data such as power and burnup distribution, refuelling rates, zonal powers, and recent fuelling history.

Approach

The program developed is called Fuel Channel Selection (FCS). This program was written by personnel at Point Lepreau Nuclear Generating Station and it involved the close participation of fuelling engineers. FCS was written in Fortran 77, and it was developed and installed on a VAX 3600. FCS is used either in the role of decision maker or assistant.

Expert VAX Operator

This project was directed at the operator of the VAX computer system.

Objective

The objective of this work was to develop an assistant to the operator of the VAX computer system. The primary emphasis was to rapidly enhance skills and reduce training time.

Approach

The approach was to develop a series of rules and procedural-based knowledge which would put the capability of experts at the disposal of relatively inexperienced operators.

Emergency Operating Procedures

Station computers control nuclear plants. Features are being added to incorporate Emergency Operating Procedures into a computerized monitoring system.

Objectives

Emergency Operating Procedures exist in nine three-ring binders. One binder has been provided for each procedure. The objective was to design a computer program to recognize Emergency Operating Procedures entry points and to provide subsequent monitoring once a specific procedure has been turned on.

Approach

Ergonomics was one of the key design considerations in the implementation of the Emergency Operating Procedures monitoring system. It has been proposed that a "gateway" computer be added to the station control computers to enable the remote acquisition of data via the Ethernet Local Area Network. The program was written in Assembler Language.

A Systems Status Monitoring System

Currently, databases of system parameters are recorded at the Point Lepreau Generating Station. By plotting the trends of these parameters, plant engineers can deduce the status of systems and systems equipment. This project will look at adding an expert system front end to the data recording and plotting programs.

Objectives

The objective of this project is to develop an expert system to monitor the status of certain power plant systems and equipment. Initially the expert system will be developed to monitor the Liquid Zone Control System, but the initial prototype will be designed so that it may be extended to include other systems.

Approach

A rule-based expert system provides a convenient format for representing the data and procedures required for a computer-based systems status monitoring system. An expert system format is designed to accommodate the numerical computations which are the basis of a conventional computer system, but also to accommodate the heuristic or practical knowledge of a human expert.

It is intended to develop a prototype using the Automated Reasoning Tool.

PRECARN

PRECARN is a consortium of more than 30 Canadian corporations drawn from a broad cross section of economic activity across Canada. New Brunswick Power is one of the corporations which make up PRECARN and which provides financial support for it.

PRECARN is involved in projects that relate to artificial intelligence and robotics. New Brunswick Power and the university have made an application to PRECARN on the subject of interactive operating manuals for power plants using an expert system.

Conclusions

A brief introduction to, and description of, expert systems applications in New Brunswick Power is given. The applications are quite diverse, and a significant number involve the Point Lepreau Nuclear Station.

For futher information contact:

Bryan K. Patterson
Point Lepreau Generating Station
New Brunswick Power Commission
Point Lepreau, New Brunswick
E0G 2H0

Eugene Hill
Department of Electrical Engineering
University of New Brunswick
Fredericton, New Brunswick
E3B 5A3

Price Waterhouse

Price Waterhouse has identified expert systems as one of its core technologies within its Information Technology consulting practice.

The firm has been active in the field of artificial intelligence, specifically knowledge-based systems, since 1988, when the Emerging Technologies Special Practice Unit was created. Over the last two years, a strong knowledge engineering capability has been established and Price Waterhouse now has expert systems consulting and development staff in Toronto, Ottawa, and Winnipeg. Our emphasis has been on the hiring of professionals who have practical implementation experience through the development of systems which are in use.

The Canadian group is complemented by similar professionals in the United Kingdom, Australia, and the United States. There is close cooperation with these groups, particularly the United States, where, in addition to AI projects, there is joint work proceeding in the fields of image processing, smart cards, and EDI.

Price Waterhouse also invests heavily in applied research and operates an Advanced Technology Center in Menlo Park, California. While the Center is tracking a range of technologies including imaging technology, large-scale document management, CASE, etc., it has a particular emphasis on AI. Over the last two years, intenationally respected researchers and educators such as Dr. Paul Turner, Xerox Parc, Dr. Richard Fykes, one of the developers of KEE, Dr. Beau Shiel, and Dr. David King, of Techknowledge, have joined the Tech Center to form the core of its research capability. The Center provides specialized consulting, education, and application research services to Price Waterhouse offices around the world and is currently developing knowledge-based systems to support specialized PW practice areas, such as Tax.

Reflecting a general market condition characterized by a slow adoption of expert system technology, early PW contracts were largely without strong vertical market focus — the firm went where the interest lay. Representative projects included:

- expert system to screen proposals for the creation of new standards
- expert system to assist in risk assessment of medical devices
- analysis of potential to use KBS in the assessment of impact of new technologies on healthcare cost containment
- expert system to support the application of competition legislation to the screening of proposed new contests
- analysis of KBS opportunities for a major broadcasting organization
- development of AI seminars for the Ontario Government
- development of AI seminars and identification of potential applications for Environment Canada
- prototype expert system to assist in the application of hazardous control point analysis to the design of chilled food manufacturing processes
- analysis of the role of AI technology in naval air warfare systems
- expert system to support a microcomputer help desk service
- expert advisory system to support chemotherapy treatment in the Manitoba Cancer Treatment and Research Foundation's outreach clinic program
- technical support to the enhancement of expert systems for the Canada Centre for Remote Sensing

While this list reflects widely diverse subject matter, it also indicates the growing penetration of expert systems interest and application. Today we are finding that client interest is stronger than ever and that the projects being pursued are significantly larger and more complex.

AI activity within Price Waterhouse will continue to be focused on knowledge-based technology and the efforts of the firm will concentrate on:

- broadening the understanding of the technology within the noninformation technology practice units of the firm, for example within the auditing practice

- continuing the development of methodology to govern KBS development and maintenance and integrating this methodology with established systems management methodology
- moving to a stronger vertical market focus whereby KBS technology will be coupled to our established specialized consulting practices, e.g., health care, environment, financial, and executive information systems
- continued development of KBS to support the large audit and tax practices
- continued emphasis on the integration of KBS into conventional information systems and the provision of KBS training to conventional systems specialists

Price Waterhouse holds the strong belief that knowledge-based systems will form an important component of all future information systems and that the use of KBS practice aids will be essential to the continued delivery and growth of accounting and consulting services in the 90's. The firm will be increasing its program of investment in AI.

For further information, please contact:

Mr. G. Thomas
Thomas International Inc.
c/o Price Waterhouse
180 Elgin Street, Suite 1100
Ottawa, Ontario
K2P 2K3
Telephone: (613) 238-8200
Fax: (613) 238-4798

Royal Bank of Canada

The Royal has been working with AI technology for several years. Preliminary work consisted of projects carried out in the Technology Planning and Development Department, to determine the impact and scope of the many technologies comprising the AI world.

Today it is slowly being adopted as an accepted technology to solve certain type of problems which, while not impossible to solve with conventional techniques, lend themselves to the techniques implemented in AI software. The work in AI has been mostly in the expert systems area. Both image and natural languages are being looked at as processing and data capacities of computers and storage media improves and becomes less costly.

Expert systems allow us to capture experienced staff members' business knowledge, then using the business data already available in electronic form, apply these rules/knowledge using inference processing techniques to suggest a recommendation or solution. Specifically, work is in progress in all the traditional areas of expert advisers, operation help desks, and risk assessment systems to help managers handle the everyday decisions that are becoming more complex each day. Recent emphasis on security issues related to financial transactions requires high performance AI software with complex pattern matching capabilities.

It is expected that inference-based technology will be embedded in standard business applications within the next couple of years.

Expert system techniques will shortly be rolled out to the traditional data processing areas, where they will be applied in those areas which lend themselves to the technology.

Research will continue in all other areas of AI.

For further information, please contact:

Elo Laugesen, Manager
Technology Projects
Royal Bank of Canada
315 Front Street West
Toronto, Ontario
M5V 3A4

Shell Canada

Abstract

Shell Canada has taken expert system technology, a subset of artificial intelligence, from a status of applied research four years ago to a proven bottom line impact technology today. This paper discusses how this was accomplish, what lessons were learned, and some future objectives.

The Beginning

Shell Canada has been working in AI technology since 1986. Initially, several projects were carried out in the Research Department to help determine the real potential of AI. Today it is being used as an accepted technology to solve a class of problems which are very difficult or impossible to solve with conventional techniques.

The work in AI has been predominantly in a subset of AI referred to as expert systems. Expert systems capture the business rules familiar to an expert in a given problem situation and use a computer inferencing technique to postulate a solution which best fits all the given rules.

In 1986 expert system technology was largely unproven in our industry. There were a few classic cases of success in specialized areas but no proof of applicability to the typical problems of a corporation like Shell. The decision to get involved in this new technology, as is always the case with any new technology, was based on the belief in a vision of how it could contribute to the well-being of the corporation.

The Vision

The vision for expert systems had three components. First, there were a class of problems which were difficult if not impossible to solve with conventional algorithmic methods using a computer, but which humans can solve readily. The earth sciences as used by petroleum companies and the operational aspects of a petroleum company have many problems which do not have exact calculable answers but which are solved by human judgements based on well-tested rules of thumb. Expert systems was a technology that promised to provide a computerized solution to this class of problem.

Second, in commercial areas of the business, there were situations where business policy was very difficult to translate into a sequence of computerized steps and where frequent policy changes required costly change to computer systems. These included areas like tax rules, royalty rules, personnel policy, and payroll policy. All these areas had been computerized, but were expensive to maintain and were subject to unreliability because of frequent changes. Expert system technology offered a better way to build the components of these systems that were based on a set of business rules rather than a set of mathematical calculations. The hope was that these systems could be built more quickly and be more reliable and less costly using this technology.

Third, there was a lot of experience and expertise in the minds of our senior people that was difficult to capture and pass on to the less experienced staff. This applied to all areas of the company from the very technical field operations, to the financial operations, and to personnel and pension policy. Expert system technology again promised some hope for capturing this knowledge and making it available 24 hours a day. The knowledge could also be used in several locations simultaneously, including remote ones.

The Process

In 1986 it was decided to dedicate resources in the research department to see if expert system technology could be incorporated into the company's operations, to learn how to go about building such expert systems, and to identify some of the difficulties of this technology. The research department provided an ideal incubator environment which freed the projects from day-to-day operational pressure so that the learning objectives could be met.

The objective of this work was to apply technology that was readily available from vendors in the form of development tools and from universities in the form of know-how and experience to Shell's specific situations. We chose to use expert systems shells as opposed to more basic languages like LISP and Prolog. The shells are like a template for building a system. They allow storage of rules in a rule base and they provide an inference engine for finding solutions which fit the rules. This means systems can be built quickly. Shells are somewhat more restrictive than LISP or Prolog but we found the restrictions were minimal for our target applications, compared to the decreased learning curve and the saving in development time.

The Result

To date we have built about a dozen systems. The majority of these are being used and providing payback, and those that are not being used taught us some very important things about how not to build expert systems. Some of the projects, which have been in many different areas, are listed below.

- *Air Drilling Advisor*: This system advises drilling engineers on the suitability of using air drilling in a given situation and on the engineering parameters.
- *Cementing Advisor*: This system advises drilling engineers on well cementing operations.
- *Credit Advisor*: This system advises on the credit risk of applicants for credit. It reduces the time to do credit evaluations and provides a consistent screening threshold.
- *Toxicology Advisor*: This system advises a toxicologist by suggesting the toxicology classification of a product based on the product composition as defined in a corporate database.
- *Help Desk Support*: Two systems were developed. They provide realtime support for a help desk operator in two different help desk situations. They broaden the expertise of a help desk operator and can be effective training tools.
- *System Maintenance Advisor*: This system advises a computer system support person about the probable cause of failure and about possible remedies and cross impacts to be aware of when making repairs.
- *Process Operator Advisor*: This system advises a process operator on recommended control action based on the trends and status of process parameters like temperature, pressure, and flow rates.
- *User Interfaces*: Several systems have been developed to advise the users of complex computer systems like reservoir models or seismic processors on how to set up a job or diagnose why it failed.

These systems are in various states of development. The majority are in operational use with positive bottom line impact. Some are still in the "proof of concept" mode where they are functional but for either technical or cultural reasons are not in production use.

In addition to the above we have a list of about 40 potential projects which could be started if resources were available. We have developed a screening process to evaluate new opportunities before committing resources. The screening process asks pertinent questions to determine if expert system technology is suitable for the given application.

Conclusion

The first two promises of expert systems have been shown to be true. We have developed systems we could not have built with conventional technology, and we have developed some applications quickly and easily which would have been very difficult with traditional techniques. The last promise of capturing expertise in a transferrable form has been demonstrated in a narrow knowledge domain and the problems of capturing a broad knowledge domain are being investigated.

There are, of course, difficulties. Knowledge acquisition, that is the process of encoding an expert's knowledge into a set of rules suitable for the expert system shell, is still an art that needs to be learned by practice. Local universities and some product vendors, however, are doing some promising research in this area. Another difficult area is cultural acceptance of this technology. Because of the implied human characteristics of an expert system, people react differently to the intrusion of such a system. More cultural and human factors engineering is required to ensure the systems will be accepted.

We are currently in the process of moving this technology into the mainstream of MIS activity so that every project can be evaluated for opportunities to benefit from expert system technology. We expect many conventional systems will have components that are rule based and therefore better and easier done with expert system techniques. There are good tools available on the market today which allow the integration of expert system technology with conventional systems development.

We confine our work to applied research aimed at real problems. Our initial work was based on basic research done by Shell's Royal Dutch laboratories in Europe in collaboration with Carnegie-Mellon University. We try to build as much as possible on the basic research done by Canadian universities, government organizations, and industry organizations like PRECARN.

The real power of computers in the 1990s will be to augment the human intellect in the same way that automobiles extend the human leg. The computer already has a large infallible memory. Expert system technology will add to this an ability to reason consistently and to do more exhaustive inferencing. This will complement humans' limited forgetful memory and subjective shortcut reasoning. We feel this computer assistance will give our knowledge worker the productive edge they will need to keep Shell strong and competitive in the 90s and beyond.

For more information, please contact:

Wally Reimer
Shell Canada Ltd.
400 4th Ave SW
Calgary, Alberta
T2P 2H5
Telephone: (403) 232-4231
Fax: (403) 232-4840

Section 6

Appendices

Appendix 1: Canadian AI Capability Survey: December 1989

Introduction

This document summarizes the results of the AI Capability in Canada study performed by Price Waterhouse for Industry, Science and Technology Canada.

Objectives of the Survey

The Technology Policy Branch of Industry, Science and Technology Canada commissioned Price Waterhouse to conduct a survey of artificial intelligence capabilities in Canada and to create an electronic database of the survey results.

The database will be maintained on a continuing basis and will be used by government departments and agencies to identify:

- organizations who could develop AI systems or AI services under contract to the federal government
- companies and universities that might participate in R&D projects
- firms that could benefit from technology transfer from government to industry
- experts who could provide advice on the scientific, technological, and business merits of proposals coming from government departments

In the interest of promoting Canadian AI capabilities, the database is available to others with similar needs. A broad spectrum of potentially interested participants was identified and included:

- computer consulting and systems integration companies
- corporations or corporate divisions which apply this technology to their business operations and/or products
- research centers
- university researchers specializing in artificial intelligence or using it to support their research

Federal government departments, agencies, and laboratories were specifically excluded from the survey. Participation in the survey was strictly voluntary.

Methodology

The mailing list of potential survey participants was compiled from the following sources:

- internal Price Waterhouse mailing lists (from seminars, personal knowledge, other historical sources)
- OASIS contractors list
- DSS contractors list
- BOSS source list
- CIAR membership list
- PRECARN membership and project lists
- NRC AI Committee members

Other potentially interest parties were identified from returned and completed survey questionnaires and from word-of-mouth. Additionally, a notice of this survey was placed in the October 1989 issue of *Canadian Artificial Intelligence*.

At the end of November, approximately 375 follow-up calls were made to organizations and individuals who were known to be in AI and who had not yet replied. Many of these sent questionnaires as a result.

Each questionnaire was scrutinized for completeness and consistency and follow-up calls were made when information was unclear or missing. Some standardization of company types (e.g., AI consulting firm, systems integration company, etc.) was performed. Categories for specializations, domains, and contracting resource types were developed (based on responses) and the questionnaire responses were allocated to these types.

It should be noted that the "Expert Systems" specialization category is very broad and includes planning, assumption-based reasoning, modeling, and simulation. "Robotics" also includes tactile systems and robotics-related sensor fusion.

All questionnaire data were entered into a dBASE III+ database system.

Summarizing the Survey Results

Response

An estimated 1373 survey questionnaires were mailed in the last two weeks of September 1989. Table 8 shows the breakdown of replies as of the end of December 1989.

Note that the number of survey participants is the number of people to whom a questionnaire was sent. The total number of questionnaires received is less than this number since people grouped themselves.

Three hundred eighty-three organizations and individuals, or 28% of those contacted, sent a reply stating that they were in the AI business or were doing AI work.

Table 8. Survey response.

Number of survey participants	383
Number of refusals to participate (from follow-up calls)	31
Number of "NOT-IN-AI" replies	90
Number of organizations/individuals on the mailing list	1373
Response rate	37%

Summarizing the Responses

The discussion in this section is based on the contents of the 312 questionnaires received before November 27, 1989. Questionnaires were completed by companies, laboratories, research centers, university departments, or individuals. When the term organization is used in the following discussion it refers to any group or individual who completed a questionnaire.

Appendices B to E of the original report contain a list of the AI specializations, domains, products, and contracting resource types which are in the database.

Most respondents supplied all the information requested. However, one company considered the entire questionnaire confidential data, while several others considered any question asking for the number and background of its employees as confidential. Some supplied numbers with the proviso that only the government would have access to the data.

In all cases, where counts were designated as confidential or labeled as unavailable, a "-1" was entered into the database. Other confidential information was not entered at all. Therefore, the results summarked in the sections below do not include confidential or unavailable information.

Questionnaires by Province

Figure 13 shows a breakdown by province of the number of organizations who responded to the survey. Figure 14 shows the number of employees, by province, represented by responding

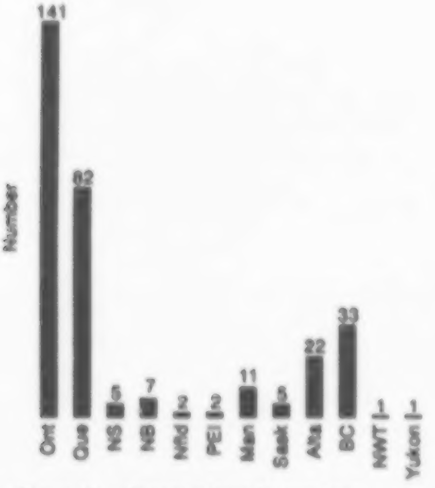

Figure 13. Organizations by province.

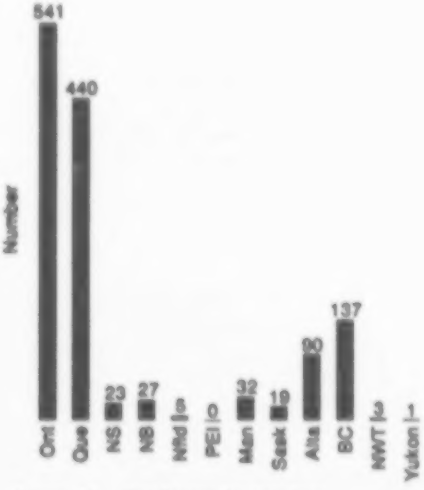

Figure 14. Employees by province.

Table 9. Total number of people represented in the survey notes.

	Full-time (Average)	Part-time (Average)
Corporate/Individual (195 questionnaires)		
Students	20(0)	64(0)
Employees	573(3)	398(2)
University/College (117 questionnaires)		
Students	604	91
Employees (including faculty)	276	71

Notes

1. The total number of part-time employees is probably high since a part-time employee could have been reported on more than one questionnaire.

2. The total number of students presented with the university data is undercounted: it was unavailable in some instances.

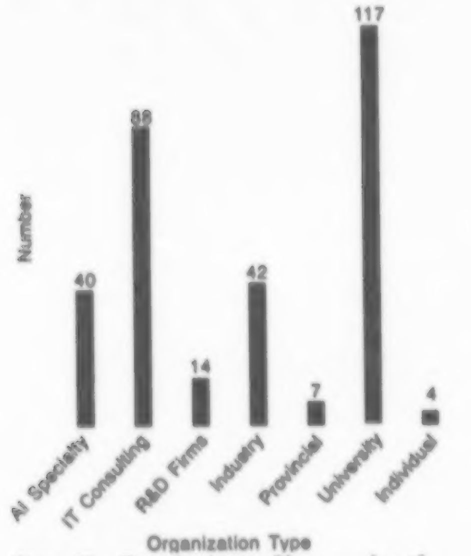

Figure 15. Organization profile — number of organizations.

organizations who are involved in AI on either a part-time or full-time basis.

People in AI

Table 9 characterizes the number of employees and students (full- and part-time) represented by survey respondents.

Organization Type

Figure 15 shows the number of organizations by major organization types.

Figures 16 and 17 show the number of full-time and part-time employees, respectively, by major organization type.

Table 10. Number of people in each AI specialization.

Specialization	Number of people
Robotics	143
Expert systems	631
Knowledge representation	23
Reasoning	2
Programming Language and Tool development	211
Knowledge acquisition	264
Intelligent databases	190
Intelligent user I/F	189
Intelligent computer-assisted instruction	177
Neural nets	126
Machine learning	93
Natural language processing	115
Machine translation	41
Speech recognition	43
Vision	154

Note:

1. Students are undercounted: university respondents tended not to be able to classify the specializations of students.

AI Specializations

A full list of the AI specializations found in the database is given as Appendix B of the original report. The number of people in each of the major specializations is given in Table 10.

Educational Background

Table 11 shows the educational background of survey respondents.

Experience

Figure 18 shows the number of people in each category of AI experience. Please note that since the wording of the question excluded AI,

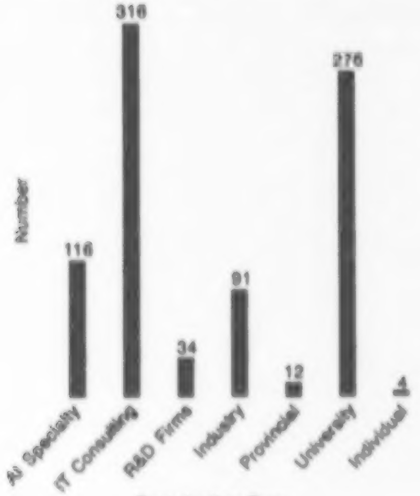

Figure 16. Organization profile — employees — full-time AI involvement.

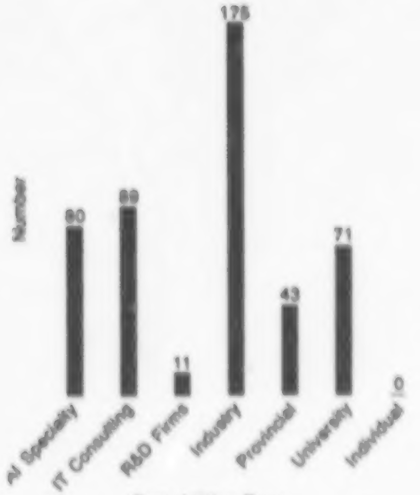

Figure 17. Organizations profile — employees — part-time AI involvement.

Table 11. Educational background of employees and students.

Educational background	Number of people	Educational background	Number of people
Bachelor in AI	164	AI courses	417
Masters in AI	308	Masters using AI	259
PhD in AI	206	PhD using AI	184

Notes:

1. The number of people with a "PhD in AI" includes PhD students, people with a PhD in AI, and professors of computer science and engineering who have sufficient experience and background to have the equivalent of a PhD in AI. The "Masters in AI" includes people with a Masters in AI and Masters students.

2. One respondent considered this question discriminatory and refused to reply. Several respondents complained that other degrees were not being tabulated.

3. Some respondents placed Masters students in the "Masters in AI" category while others may have put them in the "Bachelor in AI" category or omitted them all together. Correspondingly, PhD students may have been placed in the "Masters in AI" category or the "PhD in AI" category. Where this was obvious, it was corrected. Where it was not obvious, the data may be inconsistent

4. The "AI Courses" category was meant to refer to the number of employees who have taken AI courses but who do not have a formal degree in the area. Some respondents interpreted this as asking for the total number of courses taken. Again, where this was obvious it was corrected; however, some errors may remain.

5. All categories have some potential to include double-counting for those cases where students were reported by both the university and by their employers.

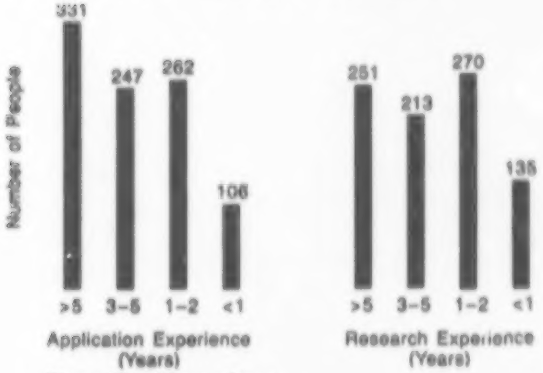

Figure 18. AI experience.

application experience figures may include people without AI experience. Where this was obvious, it was corrected. As wll, all categories have some potential to include double-counting for those cases where students were reported both by the university and their employers.

Available Resources

A full list of contracting resource types appears as Appendix E of the original report prepared for Industry, Science and Technology Canada. A summary of available personnel for each of the larger contracting resource classifications is presented in Table 12 followed by the number of

Table 12. Available resources.

	Short-term	Long-term
AI researcher	119	102
AI system designer	427	414
AI system coder	208	202
Knowledge engineer	212	168
AI strategic planner	82	71
AI project manager	13	8
DP design support for AI	93	145
AI trainer(seminar developer)	14	7
AI system integrator	2	5
Number of organizations or individuals interested in:		
Advisory work	268	
Cooperative projects	295	
Technology transfer opportunities	297	

Notes:

1. Participants frequently put one person in several categories, i.e., a person may be both a qualified knowledge engineer and an AI system designer.

2. Part-time person may be double-counted here, i.e., they may have been reported by more than one employer.

3. The "DP design support for AI" category includes some personnel not counted elsewhere: some organizations included some of their non-AI personnel here.

organizations interested in advisory positions, cooperative projects, and technology transfer opportunities.

The organization index, which appeared as Appendix F of the original report, is reproduced here as Table 13.

Table 14. Organization index.

Organization	City	Province	Contact	Telephone
222 3000 Inc.	Québec	Québec	Louis Rodrigue	418-683-4643
4th Wave Computers Ltd.	Toronto	Ontario	Mr. Peter Caven	416-752-5976
AAL Autometrics Associates Ltd.	Victoria	British Columbia	Mr. Brian H. Rowbottom	604-652-2679
Aark Consultants Ltd.	Chelsea	Québec	Dr. Gord Runtz	819-827-0054
Aastra Aerospace Inc.	Downsview	Ontario	H. Scholaert	416-736-7070
ACDS Systèmes Graphiques Inc.	Hull	Québec	M. Richard Bourgouin	819-770-9631
Acquired IntelUgence Inc.	Victoria	British Columbia	Dr. Brian Schaefer	604-383-8935
Adelphi	Gloucester	Ontario	Charles A. Rovira	613-731-1519
ADS Informatique Inc.	Québec	Québec	M. Pierre Baillargeon	418-648-0858
Aesir Management & Technologies Ltd.	Thornhill	Ontario	Mr. H.K. Thadaney	416-731-6335
AI Systems Consulting Ltd.	Calgary	Alberta	Michael P. Wingham	403-287-2351
AIT Advanced Information Technologies Corporation	Nepean	Ontario	Graeme Jones	613-226-7800
Alberta Government Telephones	Edmonton	Alberta	Vera Sanger	403-493-4706
Alberta Research Council	Calgary	Alberta	Dr. Ernest Chang	403-297-2671
Alcan Aluminium Ltee	Montréal	Québec	Vinko Potocnik	418-699-3332
AMS Management Systems Canada Inc.	Ottawa	Ontario	Mr. Martin Laplante	613-232-7400
Anaheim Technologies Inc.	Montréal	Québec	Ara Altounian	514-745-2272
Andersen Consulting	Toronto	Ontario	William M. Tatham	416-947-7773
Andrew Gullen	Toronto	Ontario	Andrew Gullen	416-767-3862
Anthony Pau & Associates Inc.	Toronto	Ontario	Anthony Pau	416-580-1566
APG Inc.	Ottawa	Ontario	Andre B. Boissonneault	613-782-2225
Applied AI Systems, Inc.	Kanata	Ontario	Mr. Takashi Gomi	613-592-3030
ARCS Labratory	Burnaby	British Columbia	Bruce Kenny	604-432-8585
Array Systems Computing Inc.	Downsview	Ontario	Linda Davis	416-736-0900
ART-EXPERT INC.	Laval	Québec	Fadel Fodda	514-622-4313
Artificial Intelligence Mgt & Dev Corp	Toronto	Ontario	Dr. Jack L. Edwards	416-488-6068
ASSEAU inc.	Québec	Québec	Francois Ouzilleau	514-272-0774
Atlantic Nuclear Services Ltd.	Fredericton	New Brunswick	Dr. C. Keith Scott	506-458-9552
Atomic Energy of Canada Ltd.	Chalk River	Ontario	Ramnik R. Shah	613-584-3311
Autonetics Research Associates Inc.	Sooke	British Columbia	Dr. Kathleen H.V. Booth	604-642-5352
Avec Technicals Service	Mississauga	Ontario	George Klima	416-855-2924
B.C. Advanced Systems Institute	Burnaby	British Columbia	Jeff Berryman	604-435-0551
Behavioral Team	Toronto	Ontario	Dr. Ben Barkow	416-656-6776
Bishop's University	Lennoxville	Québec	Marc Kaltenbach	
BMP Energy Systems Ltd.	Calgary	Alberta	Mr. Dana Krause	403-234-7979
BNH Expert Engineering Software Inc.	St. Laurent	Québec	Mr. Jihad Bahlis	514-745-4010
Brant Technologies Inc.	Mississauga	Ontario	Ross G. Hopmans	416-238-9790
British Columbia Telephone Company	Burnaby	British Columbia	Susan Lash	604-432-2399
Brock University	St. Catharine's	Ontario	Dr. James Bradford	416-688-5550
Brock University	St. Catharine's	Ontario	Dr. Jerzy A. Barchanski	
C-Team Systems Inc.	Nepean	Ontario	Mr. Don Gribble	613-727-8224
CACI Systems and Technology Ltd.	Edmonton	Alberta	Gerald Lavigne	403-471-4421
CAE Electronics Ltd.	St-Laurent	Québec	Dr. Spero Mensah	514-341-6780

Table 13. Continued.

Organization	City	Province	Contact	Telephone
Campbell & Associates	Telkwa	British Columbia	Alan N. Campbell	604-846-5853
Canadian Astronautics Limited	Ottawa	Ontario	Alexander J. M. Baillie	613-820-8280
Canadian Marconi Company Ltd.	Montréal	Québec	Mr. John Haberl	514-341-7630
Canadian National Railways	Montréal	Québec	Peter Bilissis	514-399-7003
Cancer Control Agency of British Columbia	Vancouver	British Columbia	Doug Hoffart	604-877-6020
Canpolar East Inc.	St. John's	Newfoundland	E.M. Reimer	709-722-6067
Canpolar Inc.	Toronto	Ontario	Dr. James R. Rossiter	416-487-1581
Cariboo College	Kamloops	British Columbia	Roelof Kars Brouwer	604-828-5219
Carleton University	Ottawa	Ontario	Dr. F. Oppacher	613-788-4333
Carleton University	Ottawa	Ontario	Dr. J.W. Chinneck	613-788-5733
Cavalier Information Tech. Systems Ltd.	North Gower	Ontario	Mr. George Piskor	613-489-3747
CDSL (Co-operators Data Services Ltd.)	Regina	Saskatchewan	M. Grandy Rankin	306-347-6488
CEMTECH	Ottawa	Ontario	Dr. John de Mercado	613-722-2222
Centre de Recherche Informatique de Montréal	Montréal	Québec	Susan Dube	514-848-8686
CGI	Toronto	Ontario	A.E. Thomas	416-653-0020
CIBS Corp	Willowdale	Ontario	J.H. Vowles	416-756-1730
Cogentex, Inc.	Montréal	Québec	Dr. Richard I. Kittredge	514-737-4367
Cognex	North York	Ontario	Dr. Rainer Von Konigslow	416-498-5335
CogniCom	London	Ontario	Zenon W. Pylyshyn	519-661-2063
Cognisys Consultants Inc.	Montréal	Québec	Pierre Mycholtchouk	514-599-5714
College Champlain	St. Georges de Windsor	Québec	Pierre Dion	819-564-3680
Comdale Technologies Inc.	Toronto	Ontario	Paula Sprentz	416-252-2424
Comintel Corp.	Toronto	Ontario	Donald C. Brady	416-699-1085
CompEngServ Ltd.	Ottawa	Ontario	David G. Bowen	613-563-1920
Computer Aided Technology	St. Laurent	Québec	Joseph Vybihal	514-332-4678
Computer-Based Training Systems Ltd.	Calgary	Alberta	Dr. M.G. Massey-Hicks	403-275-1680
Computing Devices of Canada	Ottawa	Ontario	Dr. T. Gajdicar	613-596-7292
Concept 5 Computer Systems Inc.	Brampton	Ontario	Mr. Edward J. Cole	416-846-4775
Concordia University	Montréal	Québec	Dr. C.Y. Suen	514-848-3000
CORECO INC.	St-Laurent	Québec	Faraj Nakhleh	514-333-1301
Crane Management Consultants Ltd.	Vancouver	British Columbia	Derek DeBiasio	604-688-1245
CRS PLUS	Burlington	Ontario	Joseph Nunn	415-639-0086
Curvet Information Systems	Oxford Mills	Ontario	David Shaw	613-258-2468
Cyberworks Inc.	Orillia	Ontario	V. Burhanpurkar	705-325-6110
Cytechnics Ltd.	Nepean	Ontario	Hazen D. Halverson	613-828-4277
D&S Knowledge Systems Inc.	Calgary	Alberta	Dave Jaques	403-268-6500
D. Stuart Conger and Associates	Ottawa	Ontario	D. Stuart Conger	613-729-4913
D.I.L. International Inc.	Sainte-Foy	Québec	Claude Labelle	418-687-9788
Dalhousie University	Halifax	Nova Scotia	Alan Gray	902-424-3632
Dalhousie University	Halifax	Nova Scotia	Dr. Jan A. Mulder	902-424-3356
Daniel Coudry LNR	Québec	Québec	Daniel Coudry	418-524-2854
Datem Ltd.	Nepean	Ontario	Mr. James Isaac	613-225-5919
Dearborn Chemical Co. Ltd.	Mississauga	Ontario	David Young	416-279-2222
Delphisoft Inc.	Sainte-Foy	Québec	M. Franck Schmitt	418-653-1765

Table 13. Continued.

Organization	City	Province	Contact	Telephone
Deltron Consulting Services Ltd.	Ottawa	Ontario	John Fodor	613-761-9700
Dewey McMillin & Associates Ltd.	Victoria	British Columbia	J.M. Dewey	604-477-5849
Diamond Software Associates	Burnaby	British Columbia	Robin McGillveray	604-439-7248
Diffracto Limited	Windsor	Ontario	Omer L. Hageniers	519-945-6373
Digital Equipment of Canada Ltd.	Toronto	Ontario	Mr. Gordon Graham	416-597-3181
Dorval Aviation Software, Inc.	Dorval	Québec	Bradley A. Longworth	514-633-9078
Dow Chemical Canada Ltd.	Sarnia	Ontario	Dr. Nicholas Darby	519-339-3922
DTI Telecom Inc.	Montréal	Québec	Daniel Y. Zann	514-845-3205
Dynacon Enterprises	Downsview	Ontario	Dr. Glen Sincarsin	415-667-0505
Eclin Corporation	Mississauga	Ontario	Denis Dorval	416-624-2348
École des Hautes Études Commerciales	Montréal	Québec	Claude Chriqui	514-343-6484
ECS - Power Systems Inc.	Ottawa	Ontario	Francois Lemay	613-236-3920
Engel & Townsend	Toronto	Ontario	Dr. M. Townsend	416-698-6376
ESSA Ltd.	Vancouver	British Columbia	Tim Webb	604-689-2912
Expert Solutions	Toronto	Ontario	Dr. Peter Davies	416-769-5695
Expertware Development Corporation	Victoria	British Columbia	Mr. Andrew D. Laks	604-381-5955
Feldgajer & Feldgajer Engineering Consultants	Waterloo	Ontario	Oleg Feldgajer	519-885-6316
Ferguson Simek Clark	Yellowknife	NWT	Garry Karst	403-920-2882
First Watercount Systems Corp.	Vancouver	British Columbia	Mr. Yves Bajard	604-876-2440
Fitzpatrick Engineering Ltd	Edmonton	Alberta	Arthur Fitzpatrick	403-438-1581
Fleet Technology	Calgary	Alberta	Dr. Alan Murray	403-253-4883
Flexpert Systems	Kemptville	Ontario	Mr. Floyd Kelly	613-258-5236
Foretell Technologies Limited	Stellarton	Nova Scotia	Mr. Vijay Sankar	902-755-1884
Francois H. Martel, M.D.	Cornwall	Ontario	Francois H. Martel	613-937-0497
G. Dentman Computer Specialists, Inc.	Leduc	Alberta	Glenn Dentman	403-986-3729
Gandalf Data Ltd.	Nepean	Ontario	Zbigniew Mikolajuk	613-723-6500
Gellman Hayward & Partners Ltd.	Toronto	Ontario	Dr. Harvey Gellman	416-862-0430
Gellman Hayward & Partners Ltd.	Calgary	Alberta	David D. Fisher	403-233-2713
Geyser Informatics Inc.	Montréal	Québec	Normand Lassonde	514-842-5719
Glen Cooper & Associates Ltd.	Vancouver	British Columbia	Glen Cooper	604-681-2054
Godin London Incorporated	London	Ontario	Mr. Gaetan Godin	519-679-8290
Golder Associates	Calgary	Alberta	Mr. Gordon L. McClymont	403-259-3413
Great-West Life Assurance Company	Winnipeg	Manitoba	David Erbach	204-946-7461
Homestead Computer Services Ltd.	Winnipeg	Manitoba	Gary Craven	204-284-9830
Honeywell Limited Advanced Technology Centre	Markham	Ontario	Dr. Blake Reid	416-479-5100
Horler Information Inc.	Ottawa	Ontario	Dr. David N.H. Horler	613-594-5155
Hydro-Québec	Varennes	Québec	Gilles Desrochers	514-652-8321
Idon Corporation	Ottawa	Ontario	Herbert G. Bown	613-722-8101
IEA Consulting Group Ltd.	Charlottetown	P.E.I.	George W. Mason	902-368-5014
Imago Manufacturing Ltd.	Ottawa	Ontario	Roy Ball	613-226-7890
Infoknowledge Inc.	Orleans	Ontario	Andre Hintis	613-830-2720
Information Fusion Systems Corporation	Nepean	Ontario	Dr. Gregg Van Volkenburgh	613-726-0934
Innovus Inc.	Hamilton	Ontario	Robert S. Sciuk	416-529-8117
Intellectia	Montréal Nord	Québec	Houda Araj	514-325-7135

Table 13. Continued.

Organization	City	Province	Contact	Telephone
Jasco Research Ltd.	Sidney	British Columbia	Dr. J.A. Scrimger	604-656-0656
KBA - Knowledge Base Associates	Victoria	British Columbia	Marie Curtis	604-384-6433
KIS Information Systems Inc.	Vancouver	British Columbia	Carl Rogger	604-222-2252
Knoman Consulting	Ottawa	Ontario	Mark Wright	613-594-5295
Knowledge Transfer	Ottawa	Ontario	Kay James	613-748-9650
FirstMark Technologies Ltd.	Ottawa	Ontario	Mr. Barry de Ville	613-723-8020
Lafarge Canada Inc.	Woodstock	Ontario	Sam Fujimoto	519-423-6241
Lamb Cargate Industries Ltd.	New Westminster	British Columbia	Ian C. Campbell	604-521-8821
Le groupe Micro-Intel Inc	Montréal	Québec	Gilbert Paquette	514-255-6048
Les Conseillers Masycom Consulting Inc.	Gatineau	Québec	Norma Welch	819-561-2354
Les Systemes Expertex inc.	Montréal	Québec	Mario Robert	514-844-5157
LGS Group Inc.	Montréal	Québec	Jean-Francois Grou	514-861-2673
Lobbe Technologies Ltd.	Regina	Saskatchewan	Dr. A.J. Szladow	306-586-9400
Loewen Translations	Nepean	Ontario	Victor Loewen	613-723-9121
Logicware International	Mississauga	Ontario	Walter Berndl	416-629-8801
London Life Information Services	London	Ontario	Terry Wood	519-432-5281
LSG Ingenieurs-Conseil en Logiciel Enr.	St-Jean-sur-Richelieu	Québec	Richard Lemire	514-346-3762
Luc Michaud	Montréal	Québec	Luc Michaud	514-334-3268
MacDonald Dettwiler & Associates, Ltd.	Richmond	British Columbia	Timothy Bult	604-278-3411
Machina Sapiens	Montréal	Quebec	Claude Coulombe	514-733-1095
Maison de Technologies Avancees	Greenfield Park	Québec	Saadeldin El-Sedfy	514-445-7363
Manitoba Data Services	Winnipeg	Manitoba	Sharon D. Satz	204-946-6681
Manitoba Hydro	Winnipeg	Manitoba	Laura Penner, P. Eng.	204-474-3054
Marc Garneau Collegiate Ins.	Don Mills	Ontario	John D. Craig	416-429-6166
McCarney Technologies Inc.	Vancouver	British Columbia	Stefan W. Joseph	604-669-7480
McCartney Consultants Ltd.	Toronto	Ontario	M. McCartney	416-763-3102
McGill University	Montréal	Québec	Margaret Dalziel	514-398-6319
McGill University	Montréal	Québec	Thomas R. Shultz	514-398-6139
McMaster University	Hamilton	Ontario	H.A. El Maraghy	416-525-9140
McMaster University	Hamilton	Ontario	Dr. Gilles G. Patry	416-525-9140
McMaster University	Hamilton	Ontario	Ali-Reza Montazemi	416-525-9140
McMaster University	Hamilton	Ontario	Dr. S. Poehlman	416-525-9140
McMaster University	Hamilton	Ontario	Norman P. Archer	416-525-9140
McMaster University Medical Centre	Hamilton	Ontario	Karl Langton	416-525-9140
Memorial University of Newfoundland	St. John's	Newfoundland	J.M. Foltz	709-737-8628
Metropolitan Life Insurance Company	Ottawa	Ontario	Heather Swayne	613-738-5813
MicroLight Graphics Ltd.	North Vancouver	British Columbia	John Whalen	604-980-5353
Microtec Ltd.	Fredericton	New Brunswick	Dr. C. Keith Scott	506-458-9553
Microtel Pacific Research Limited	Burnaby	British Columbia	Dr. Rizwan Kheraj	604-293-6009
Ministry of Revenue	Oshawa	Ontario	Thomas Boyd	416-433-6840
MONAC	Winnipeg	Manitoba	John Craig Muller	204-943-6797
MPB Technologies Inc.	Dorval	Québec	Dr. M.P. Bachynski	514-683-1490
Netron Inc.	Downsview	Ontario	Paul Bassett	416-636-8333
Neural Systems Incorporated	Vancouver	British Columbia	Mr. Gary Josin	604-263-3667

Table 13. Conitinued.

Organization	City	Province	Contact	Telephone
New Brunswick Electric Power Commission	Fredericton	New Brunswick	Gene Maillet	506-458-3628
Nidak Associates Inc.	Toronto	Ontario	Mr. Nigel Stokes	416-861-9133
Noranda Technology Centre	Pointe Claire	Québec	Robert Spring	514-630-9521
Norman Clark	Leamington	Ontario	Norman Clarke	519-322-1387
Northern Dynamics Ltd.	North York	Ontario	Mr. John Finch	416-492-8240
Northern Telecom	Toronto	Ontario	L. Baxter	416-598-0196
Nova Scotia Power Corp.	Halifax	Nova Scotia	E.G. Achorn	
NUTAT Technologies Inc.	Ottawa	Ontario	Tim Boreham	613-235-7654
Oerlikon Aerospace Inc.	St-Jean-sur-Richelieu	Québec	Michel Rodrigue	514-358-2000
Ontario Hydro	Toronto	Ontario	H.S. Irvine	416-592-5539
Paul Egan	Charlottetown	P.E.I.	Paul Egan	902-892-2140
Pedestal Robotics	Montréal	Québec	Howard Schneider	514-931-3365
Pierre Michaud, Consultant en Informatique	Outremont	Québec	Pierre Michaud	514-276-6786
Pratt & Whitney Canada Inc.	Longueuil	Québec	Jacques Chelin	514-677-9411
Precarn Associates Inc.	Nepean	Ontario	Mr. Jean-Claude Gavrel	613-727-9576
Price Waterhouse	Ottawa	Ontario	Oliver Kent	613-238-8200
Procter & Gamble Canada Inc.	Hamilton	Ontario	Ralph J. van Engelen	416-545-1121
Prof. Michael Jenkin	Downsview	Ontario	Mike Jenkins	416-736-5053
PROGICIELS Bourbeau Pinard inc.	Montréal	Québec	Mr. Laurent Bourbeau	514-276-5382
QMC	Ottawa	Ontario	Hans Witte	613-233-3089
Queen's University	Kingston	Ontario	M.A. Jenkin	613-545-6073
Queen's University	Kingston	Ontario	Brian Surgenor	613-545-2568
Queen's University	Kingston	Ontario	Randy Ellis	
QuesTech North American Ltd.	Manotick	Ontario	Keith Migdal	613-692-0885
Quilchena Consulting Ltd.	Victoria	British Columbia	Mr. Anthony K. Copping	604-721-2205
Quinte Learning Centre	Belleville	Ontario	Michael Maloney	613-966-5603
R. Purdy Associates	Winnipeg	Manitoba	Robert D. Purdy	204-488-3651
Range-Imaging Vision Systems Inc.		British Columbia	Curt Lang	604-986-7799
Red River Community College	Winnipeg	Manitoba	Mr. Alan H. Robbins	204-632-2429
Rhodes and Associates, Human Factors	Willowdale	Ontario	Dr. Wayne Rhodes	416-494-2816
Rockwood Informatics	Kanata	Ontario	Philip J. Vincent	613-592-6726
Ron Gershon	Willowdale	Ontario	Dr. Ron Gershon	416-663-4048
RSI Robotic Systems International Ltd.	Sidney	British Columbia	Jack Wilson	604-656-0101
RTS Real Time Systems Inc.	Toronto	Ontario	Mr. Frank Simtob	416-449-5931
Ryerson Polytechnical Institute	Toronto	Ontario	Mr. R. Greer Lavery, P. Eng.	416-979-5063
S.L. Ross Environmental Research Ltd.	Ottawa	Ontario	Randy Belore	613-232-1564
Sanctuary Woods Inc.	Dunrobin	Ontario	Brian Beninger	613-832-3104
SED Systems Inc.	Saskatoon	Saskatchewan	Eldon Wig	306-933-1506
Seimac Limited and Oceanroutes Canada Inc.	Dartmouth	Nova Scotia	Geoffrey P. Allen	902-468-3007
Semantax Corp.	Montréal	Québec	Conrad F. Sabourin	514-731-4266
Senes Consultants Ltd.	Toronto	Ontario	Bruce Powers	416-764-9380
Serdek Automated Systems Inc.	Chatham	Ontario	Ron Dekker	519-351-4230
Serdula Systems Ltd.	Deep River	Ontario	Kenneth Serdula	613-584-4588
Shell Canada	Calgary	Alberta	Wally Reimer	403-232-4231

Table 13. Continued.

Organization	City	Province	Contact	Telephone
Sierra Systems Consultants Inc.	Vancouver	British Columbia	Don Schellenberg	604-688-1371
Silico Computer Systems Ltd.	Nepean	Ontario	Gerald J. Scarcella	613-727-0465
Simon Fraser University	Burnaby	British Columbia	Dr. Nick Cercone	604-291-3229
Simon Fraser University	Burnaby	British Columbia	Dr. Nick Cercone	604-291-3229
Simon Fraser University	Burnaby	British Columbia	Dr. John Dewey Jones	604-291-4679
Simulations Canada	Bridgewater	Nova Scotia	Mr. Stephen Newberg	902-688-2808
Simulcalc Inc.	Trois-Rivières	Québec	Denis Carrier	819-372-9617
Softwords, Div. of Press Porcepic Limited	Victoria	British Columbia	Ms. Ellen Godfrey	604-727-6522
Spar Aerospace Limited	Weston	Ontario	Dr. I.H. Rowe	416-856-4552
SPS	Ottawa	Ontario	W. Caswell	613-238-1700
Sussex Informatics Limited	Ottawa	Ontario	Mr. Brian Hierlihy	613-230-5444
System Builders	Ottawa	Ontario	Peter Fisk	613-230-4067
Systemes Metavision	Montréal	Québec	Joseph Heath	514-284-2840
SystemHouse	Ottawa	Ontario	Paul Renaud	613-236-9734
Systems Interface Inc.	Nepean	Ontario	B. Dowdall	416-727-5001
TDK Consulting Services	Kitchener	Ontario	Tim Kuehn	519-741-3623
Technical Security Service (TSS)	Winnipeg	Manitoba	Tony K. Mirza	204-944-8181
Technical University of Nova Scotia	Halifax	Nova Scotia	Dr. J. Barzilai	902-420-7573
Techno Scientific Inc.	Woodbridge	Ontario	Dr. Mirek Macecek	416-851-9958
Technogram Inc.	Chicoutimi	Québec	Jeffrey R. Machan	418-696-2164
Teknica Resource Development Ltd.	Calgary	Alberta	Don Schieman	403-269-4386
Tektrend International Inc.	Ville St. Laurent	Québec	D. Robert Hay	514-333-7371
The Herongate Group Inc.	London	Ontario	Susan M. Gardner	519-472-1855
The Laser Center Incorporated	Calgary	Alberta	Joyce Hodgson	403-251-6804
The Ontario Institute for Studies in Education	Toronto	Ontario	George Tracz	416-923-6641
Theos Philos Teknikos Ltd.	Scarborough	Ontario	Michael J. Sutton	416-694-0167
Thomas International Inc.	Ottawa	Ontario	Grant Thomas	613-749-4873
TransAlta Utilities Corporation	Ca'gary	Alberta	Mike Blakely	403-267-7582
Trent University	Peterborough	Ontario	Dr. S. Regoczei	705-748-1495
TRI Telepresence Research Inc., International Telepresence	Vancouver	British Columbia	Tom Mitchell	604-873-3300
Unitran Systems Inc.	Nepean	Ontario	Alan Cranshaw	613-228-0858
Université d'Ottawa	Ottawa	Ontario	George M. White	613-564-2940
Université d'Ottawa	Ottawa	Ontario	Dr. Dan Ionescu	613-564-2211
Université de Moncton	Moncton	New Brunswick	M. Rejean Hall, ing.	506-858-4307
Université de Montréal	Montréal	Québec	Patrick Cavanagh	514-343-6154
Université de Montréal	Montréal	Québec	Dr. Jean G. Vaucher	514-343-7092
Université de Montréal	Montréal	Québec	Dr. Claude Frasson	514-343-7019
Université de Montréal	Montréal	Québec	Dr. Jan Gecsei	514-343-6509
Université de Montréal	Montréal	Québec	Dr. Jean-Charles Bernard	514-340-4723
Université de Montréal	Montréal	Québec	Michel Boyer	514 343-6804
Université de Montréal	Montréal	Québec	Daniel Poulin	514-343-2136
Université de Montréal	Montréal	Québec	Guy Lapalme	514-343-7220
Université de Montréal	Montréal	Québec	Bernard Lefebvre	514-343-7477

Table 13. Continued.

Organization	City	Province	Contact	Telephone
Université de Montréal	Montréal	Québec	Daniele Archambault	514-343-7041
Université de Montréal	Montréal	Québec	George W. Baylor	514-434-7633
Université de Montréal	Montréal	Québec	Paul Bratley	514-343-7478
Université de Montréal	Montréal	Québec	Aude Dufresne	514-343-7303
Université de Montréal	Montréal	Québec	Luc Giroux	514-343-7303
Université de Montréal	Montréal	Québec	Serge Larochelle	514-343-7633
Université de Montréal	Montréal	Québec	Gisele Lemoyne	514-343-7889
Université de Montréal	Montréal	Québec	Jean-Yves Morin	514-343-7931
Université de Montréal	Montréal	Québec	Regine Pierre	514-343-7246
Université de Montréal	Montréal	Québec	Neil Frederick Stewart	514-343-7483
Université de Montréal	Montréal	Québec	Francois Major	514-343-7508
Université du Québec à Montréal	Montréal	Québec	Claude Thomasset	514-987-8397
Université du Québec à Trois-Rivières	Trois-Rivières	Québec	Sylvain Delisle	613-564-5420
Université Laval	Sainte-Foy	Québec	Dr. Denis Poussart	418-656-3554
Université Laval	Sainte-Foy	Québec	Prof. Bernard Moulin	418-656-5580
Université Laval	Sainte-Foy	Québec	Jean Roy	418-656-5472
Université Laval	Québec	Québec	Laurent Gauthier	418-656-2662
Université Laval	Sainte-Foy	Québec	Madame Denise Deshaies	613-656-5506
Université Laval	Sainte-Foy	Québec	Pierre Maranda	418-656-7687
Université Laval	Sainte-Foy	Québec	Gilles Lemire	418-656-3651
Université Laval	Sainte-Foy	Québec	Jean Serodes	418-656-2868
Université Laval	Sainte-Foy	Québec	Stanislaw Tarasiewicz	418-656-2106
Université Laval	Sainte-Foy	Québec	Diane Huot	418-656-7468
Université Laval	Sainte-Foy	Québec	Bernard R. Hodgson	418-656-2975
Université Laval	Sainte-Foy	Québec	Henri H. Arsenault	418-656-2650
Université Laval	Sainte-Foy	Québec	Frederic Flament	418-656-2431
Université Laval	Sainte-Foy	Québec	Andre Turmel	418-656-3074
Université Laval	Sainte-Foy	Québec	Jacques Imbeau	418-656-3100
Université Laval	Sainte-Foy	Québec	Francois Bergeron	418-656-7940
Université Laval	Sainte-Foy	Québec	Jean-Marc Suret	418-656-7134
University of Alberta	Edmonton	Alberta	Dr. Renee Elio	
University of Alberta	Edmonton	Alberta	Dr. Tony A. Marsland	403-492-3971
University of Alberta	Edmonton	Alberta	Leonard Cheng	403-476-9766
University of Alberta Hospitals	Edmonton	Alberta	Dr. D. Menon	403-492-6843
University of British Columbia	Vancouver	British Columbia	Dr. D.G. Lowe	604-228-3170
University of British Columbia	Vancouver	British Columbia	Daphne Gelbar	604-228-5028
University of British Columbia	Vancouver	British Columbia	Dr. Samuel C. Coval	604-228-6481
University of British Columbia	Vancouver	British Columbia	Dr. Richard S. Rosenberg	604-228-4142
University of British Columbia	Vancouver	British Columbia	Dr. Frank Navin	604-228-3158
University of British Columbia	Vancouver	British Columbia	Dr. S.F. Stiemer	604-228-6301
University of British Columbia	Vancouver	British Columbia	Robert C. Goldstein	604-224-8389
University of Calgary	Calgary	Alberta	Dr. Ian H. Witten	403-220-6780
University of Calgary	Calgary	Alberta	Dr. Edward J. Krakiwsky	403-220-7878
University of Calgary	Calgary	Alberta	David R. Hill	403-282-6481
University of Guelph	Guelph	Ontario	Dr. Mary McLeish	519-824-4120

Table 13. Concluded.

Organization	City	Province	Contact	Telephone
University of Guelph	Guelph	Ontario	David Chiu	519-824-4120
University of Manitoba	Winnipeg	Manitoba	Dr. Simonovic P. Slobodan	204-474-6443
University of Manitoba	Winnipeg	Manitoba	Dr. Mark Evans	204-474-8313
University of New Brunswick	Fredericton	New Brunswick	Dr. Bradford G. Nickerson	506-453-4566
University of New Brunswick	Fredericton	New Brunswick	Christopher Bayly	506-453-4776
University of New Brunswick	Fredericton	New Brunswick	Evelyn W. Richards	506-453-3533
University of New Brunswick	Fredericton	New Brunswick	Dr. John Christian, Ph.D., P.Eng.	506-453-4521
University of Ottawa	Ottawa	Ontario	Dr. S.J. Matwin	613-564-5069
University of Ottawa	Ottawa	Ontario	Dr. Douglas R. Skuce	613-564-5418
University of Ottawa	Ottawa	Ontario	Dr. Stan Szpakowicz	613-564-2450
University of Ottawa	Ottawa	Ontario	Dr. Tuncer Oren	613-564-5068
University of Ottawa	Ottawa	Ontario	Dr. E. Petriu, Ph.D., P.Eng.	613-564-2497
University of Saskatchewan	Saskatoon	Saskatchewan	Dr. M. M. Gupta	306-966-5451
University of Saskatchewan	Saskatoon	Saskatchewan	Dr. G. McCalla	306-966-8655
University of Toronto	Toronto	Ontario	Dr. J.K. Tsotsos	416-978-3619
University of Toronto	Toronto	Ontario	Dr. John Mylopoulos	416-978-5180
University of Toronto	Toronto	Ontario	Dr. R. Reiter	416-978-6324
University of Toronto	Toronto	Ontario	Hector Levesque	416-978-3618
University of Toronto	Toronto	Ontario	Dr. Graeme Hirst	416-978-8747
University of Toronto	Toronto	Ontario	Dr. Russell Greiner	416-978-6277
University of Toronto	Toronto	Ontario	Dr. E. Milios	416-978-6114
University of Toronto	Toronto	Ontario	Dr. I. Burhan Turksen, Ph.D., P.Eng.	416-978-8214
University of Toronto	Downsview	Ontario	G.M.T. D'Eleuterio	416-667-7707
University of Victoria	Victoria	British Columbia	Kin Fun Li	604-721-8683
University of Waterloo	Waterloo	Ontario	Dr. Robin Cohen	519-888-4457
University of Waterloo	Waterloo	Ontario	Ed Jernigan	519-888-4646
University of Waterloo	Waterloo	Ontario	Lindsay Patten	519-885-1211
University of Waterloo	Waterloo	Ontario	Dr. Grantham Pang	519-885-1211
University of Waterloo	Waterloo	Ontario	Sanjeev Bedi	519-885-1211
University of Western Ontario	London	Ontario	Dr. Zenon W. Pylyshyn	519-661-2063
University of Western Ontario	London	Ontario	Dr. Martin J. Stillman	519-661-3821
University of Western Ontario	London	Ontario	Dr. I. Gargantini	519-661-3563
University of Windsor	Windsor	Ontario	David D. Reynolds	519-253-4232
Urbanprobe Associates Limited	Toronto	Ontario	G. Bauz	416-963-8351
UTL Canada Inc.	Kanata	Ontario	Mr. Edward S. Richmond	613-591-3474
Vadeko International Inc.	Mississauga	Ontario	Dr. Peter Wright	416-821-3222
Valcom Ltd.	Ottawa	Ontario	Eric Alain	613-594-4563
VCS Technologies Inc.	Calgary	Alberta	Dan Paslawski	403-234-8711
W. James Porcher Consulting Ltd.	Toronto	Ontario	Mr. W. James Porcher	416-947-0108
W.E. Thorp Associates Limited	Ottawa	Ontario	W. Thorp	613-729-6978
Wastart Limited	Guelph	Ontario	Larry Whittle	519-763-6443
Wayne Powell	Delta	British Columbia	Wayne Powell	604-583-3048
Xerox Canada Inc.	North York	Ontario	Bryan M. Kramer	416-733-6377
Xerox Research Center of Canada	Mississauga	Ontario	Michael A. Hopper	416-823-7091
Yukon College	Whitehorse	Yukon	Glenn Piwowar	403-668-5509

Appendix 2: Canadian Society for Computational Studies of Intelligence

CSCSI/SCEIO is the Canadian Society for Computational Studies of Intelligence/Société canadienne pour l'étude de l'intelligence par ordinateur. Founded in 1973, it originally consisted of a small group of academics and graduate students interested in AI, but the society has grown enormously and now includes members from industry, business, and government.

There are currently about 800 members, mostly from Canada, but non-Canadian members are also welcome. In fact, we are now arranging regular exchange of communications between *Canadian Artificial Intelligence* and a European equivalent.

The CSCSI exists as a special interest group of CIPS for the promotion of interest and activity in AI. Under this banner, CSCSI conducts workshops and a national conference every two years. It also publishes a quarterly magazine, *Canadian Artificial Intelligence*, and sponsors the journal *Computational Intelligence.*

The 8th CSCSI/SCEIO conference was held at the University of Ottawa, May 22–25, 1990, with a central stream of technical presentations, including distinguished guest speakers in addition to submitted papers. Plans for the conference also included a one-day set of tutorial presentations, and a full program of exhibits. In previous years, conferences have been held in Edmonton, Alberta (1988), Montreal, Quebec (1986), and London, Ontario (1984).

A major benefit of membership in CSCSI is the *Canadian Artificial Intelligence*. It was started in 1974 as the *CSCSI/SCBIO Newsletter* and was reborn in 1984 through the efforts of Professor Graeme Hirst of the University of Toronto. Recent issues have been produced in Calgary with assistance from the Alberta Research Council. The magazine contains articles of general interest, descriptions of current research and courses, reports of recent conferences and workshops, news about AI companies and products, paid advertising, and announcements of forthcoming activities.

Canadian Artificial Intelligence also contains, as a regular feature, the "World Watch on AI Applications and Developments" which is prepared by the Secretariat of the NRC Associate Committee on Artificial Intelligence.

For further information concerning the CSCSI/SCEIO contact:

CIPS National Office
5th Floor, 243 College Street
Toronto, Ontario
M5T 2Y1
Telephone: (416) 593-4040

Appendix 3: Computational Intelligence

Editors: N. Cercone, G. McCalla

Associate Editors: L. Bolc, V. Dahl, K. Fuchi, B. Grosz, A.K. Mackworth, J. Mylopoulos, A. Ortony, R. Perrault, Z. Pylyshyn, R. Reit, L. Rendell, E. Sandewall, L. Schubert, A. Sloman, N.S. Sridharan, J. Tsotsos, B.L. Webber, D. Wilkins, R. Woodham

Computational Intelligence, the official journal of the Canadian Society for Computational Studies of Intelligence, is a quarterly journal first published in 1985 by the National Research Council of Canada. It contains high-quality theoretical and experimental research papers in computational (artificial) intelligence, by encouraging contributions from the following fields: knowledge representation; natural language understanding; computational vision; applications of artificial intelligence; logic programming; theorem proving; learning; cognitive science; problem solving and planning; languages and tools for artificial intelligence; speech understanding; game playing; philosophical implications; and foundations of artificial intelligence. Three special issues were published in 1988: "Taking Issue: an inquiry into Computer Understanding" (February), "AI in France" (May), and "Planning" (November). The journal is international in content and distribution and is quickly becoming one of the leading AI journals in the world.

Computational Intelligence is published quarterly. Subscription rates for 1990 are $48.00 for individuals and $120.00 for Institutions.

Appendix 4: National Research Council, Associate Committee on Artificial Intelligence

The NRC Associate Committee on Artificial Intelligence (AI) was established in April 1987 to act as a focal point and forum in Canada for issues of national concern to the scientific and engineering communities.

NRC Associate Committees of this nature have been used for similar purposes on various topics for many years. Currently other NRC Associate Committees exist for a wide variety of topics including astronomy, biotechnology, forest fire management, toxicology, and space research.

Artificial intelligence is a branch of computer science that concerns the ability of computers to perform intelligent tasks, such as those requiring recognition, reasoning, and learning. Briefly stated, AI research and applications seek to incorporate human skills and expertise into computer systems. Applications of AI are considered by many to be virtually limitless but include:

- natural language processing
- machine translation
- expert systems, such as those used in the diagnosis and treatment of disease, monitoring and diagnosing faults in equipment, industrial plants, telecommunications, and electric utility systems
- operations scheduling and process planning in manufacturing
- robot and other autonomous systems with vision or other sensing capabilities that permit them to respond to changes in their external environment

Terms of Reference

The terms of reference of the committee are:

1. To bring together specialists from various disciplines involved in the field of artificial intelligence with emphasis on business, industrial, and social applications of benefit to Canada, in order to formulate guidelines for research and development.
2. To provide advice on possible arrangements for coordinating relevant research and development in universities, government laboratories, and industry, with the aim of establishing a strong Canadian-based expertise.
3. To advise the National Research Council, federal departments, and other organizations regarding research and development requirements in the field of artificial intelligence.
4. To establish subcommittees for the study of specific problems as may be required.
5. To provide a focus for the maintenance and dissemination of technical information nationally and internationally, in the field of artificial intelligence, particularly where this will assist and strengthen existing technical societies and organizations who might ultimately continue such activity with their own resources.
6. To carry out such other activities as may be assigned by the National Research Council from time to time.

Membership

Members of the NRC ACAI serve without remuneration or honorarium, but are reimbursed at standard government travel rates for costs incurred while on committee business or attending committee meetings. The committee normally meets four times per year. Membership has been selected to provide a wide range of representation and interests including industry, academic, and government from all regions of Canada. The membership includes individuals associated with as many other bodies and organizations with AI interests as possible in order that it may avoid duplication and act as a supplement to the work of others.

Key Areas of Activity

The committee has developed a working paper on potential plans and activities. Key areas selected for priority include:

- Encouraging closer collaboration between AI users and the AI research community in university and government laboratories
- Providing advice to the federal government on strategic planning for AI and other information technologies

- Providing advice to the National Research Council in the development of its application oriented R&D program in computer technology and artificial intelligence

Membership of the Associate Committee on Artificial Intelligence 1989–1990

Mr. L. Bolger
Shell Canada Ltd.
Calgary, Alberta

Dr. Marie-Michèle Boulet
Université Laval
Québec, Québec

Dr. N. Cercone
Simon Fraser University
Burnaby, British Columbia

Dr. E. Chang
Alberta Research Council
Calgary, Alberta

Dr. R. De Mori (Past Chairman)
Centre de recherche informatique de Montréal (CRIM)
Montreal, Québec

Mr. P.G. George
MacDonald Dettwiler and Associates Ltd.
Richmond, British Columbia

Mr. K. Hayes
Canadian Labour Congress
Ottawa, Ontario

Dr. E.G. Manning
University of Victoria
Victoria, British Columbia

M. R. Manoliu
Hydro-Québec
Montreal, Québec

Dr. G. McCalla
University of Saskatchewan
Saskatoon, Saskatchewan

Dr. G.M. MacNabb
PRECARN Associates Inc.
Nepean, Ontario

Dr. J. Mulder
Dalhousie University
Halifax, Nova Scotia

Dr. R. Peacocke
Bell-Northern Research
Ottawa, Ontario

Prof. Z. Pylyshyn
The University of Western Ontario
London, Ontario

Mr. J. Scrimgeour (Secretary)
National Research Council
Ottawa, Ontario

Mr. G.F. Sekely
Canadian Pacific
Toronto, Ontario

Dr. B.J. Smith
Acquired Intelligence
Victoria, British Columbia

Mr. G. Thomas (Chairman)
Price Waterhouse
Ottawa, Ontario

Additional contributions to the report are acknowledged from:

Ms. C. Baginski
Canadian Pacific
Toronto, Ontario

Mr. David, G. Bower
CompEngServ Ltd.
Ottawa, Ontario

Dr. A.E. Collin
Institute for Robotics and Intelligent Systems
Ottawa, Ontario

Mr. V. DiCiccio
University of Waterloo
Waterloo, Ontario

Mr. Marcel Drouin
Canadian Workplace Automation Research Centre
Laval, Quebec

Mr. A. Dubois
Industry, Science and Technology Canada
Ottawa, Ontario

Dr. David G. Goodenough
Energy, Mines and Resources Canada
Ottawa, Ontario

Mr. Gordon Graham
Digital Equipment of Canada Ltd.
Willowdale, Ontario

Ms. Ann Griffith
Applied AI Systems, Inc
Kanata, Ontario

Mr. Eugene Hill
University of New Brunswick
Fredericton, New Brunsick

Dr. Zachary Jacobson
Government Consulting Group
Ottawa, Ontario

Mr. Elo Laugesen
Royal Bank of Canada
Toronto, Ontario

Mr. Andr) Lemire
Hydro-Qu)bec
Montreal, Quebec

Mr. R.T. Lewis
Transport Canada
Ottawa, Ontario

Dr. S. MacLean
Canadian Space Agency
Ottawa, Ontario

Mr. S.A. Mayman
National Research Council
Ottawa, Ontario

Mr. Bryan K. Patterson
New Brunswick Power Commission
Point Lepreau, New Brunswick

Dr. Dorothy Phillips
Communications Research Centre
Department of Communications
Ottawa, Ontario

Mr. Wally Reimer
Shell Canada Ltd.
Calgary, Albert

Mr. Vincent Taylor
National Defence Headquarters
Department of National Defence
Ottawa, Ontario

Ms. J. Walden
Natural Sciences and Engineering Research
Council of Canada (NSERC)
Ottawa, Ontario

Appendix 5: References and Related Reports

1. *Proceedings of the Workshop on the Application of Expert Systems to Transportation*. Research and Development Directorate, Transport Canada, Ottawa. December 1985.

2. *Expert Systems: Their Application in the Canadian Transportation Sector*. TP 7328E. Prepared by Cognos Advanced Technology and Applied AI Systems for the Director General, Research and Development, Transport Canada, Ottawa. January 1986.

3. *Machine Translation and Natural Language Processing: Opportunities for Artificial Intelligence in Canada*. Prepared by Cognos, February 1985. Available from the Department of Communications and the Department of the Secretary of State, Ottawa.

4. *L'Intelligence artificielle au Québec — Références utiles*. Direction générale des politiques. Ministère des communications, Gouvernement du Québec, Québec. Mai 1986

5. *Point de vue québecois sur l'intelligence artificielle*. Richard Parent pour le Ministère des communications du Québec, Québec. 1984. ISBN 2-550-11473-6

6. *Opportunities for the Use of Artificial Intelligence Systems within the Government of Canada*. Nordicity Group and Canadian Artificial Intelligence Products Corp. for the Ministry of State for Science and Technology, Ottawa. August 1986

7. Scrimgeour, J. "Robotics." *The Canadian Encyclopedia - Second Edition*. Hurtig Publishers Limited, Edmonton, Alberta 1988 p. 1881.

8. Goldenberg, A. *An Index and Guide to Robotics Research & Development in Canada*. Division of Electrical Engineering report ERB-978, National Research Council Canada. NRCC 24470 (English), NRCC 24471 (French), Ottawa, Ontario

Section 7

Index

Index

ACAI
  See Associate Committee on Artificial Intelligence
ACQUIRE 139–140
Acquired Intelligence Inc 11, 63, 139–141
Adamovits, P.J. 64–65
Advanced Process Analysis and Control System 11, 147–148
Advanced Technology Directorate 6, 66–68
Advisory systems
  See also Applications, Expert systems, Knowledge-based systems
  air drilling 175
  analysis of LANDSAT data 7, 73
  assessment of gold mill effluent plant operation 54
  chemotherapy treatment 13, 171
  credit risk assessment 14, 175
  disk failure recovery 151
  fleet operations 143
  general applications 88
  geotechnical survey planning 83
  intelligent 8, 87
  maintenance 8, 12, 89
  mapping 77
  mining 8, 91
  process operator 175
  rail car scheduling 12
  statistical mapping 7
  system maintenance 175
  toxicology 175
  weather forecasting 32
  well cementing 175
Aerodynamic models 148
AI Fellowship Program 12, 155
AI R&D fund 5
AIR Program 3, 25
Air traffic control 12, 96–97, 147, 149
Alberta Government Telephones 92, 94
Alberta Research Council 8, 41, 91–92, 95, 118, 122, 165, 192
Alcan International Limited 41–42
Alleliunas, R. 124
AMETHYST 66
APACS
  See Advanced Process Analysis and Control System
Applications
  accounting 172
  See also Advisory systems, Diagnostic systems, Expert systems, Knowledge-based systems, Monitoring systems, Planning systems, Vision systems
  aircraft maintenance 8, 10, 87, 94, 118
  biology 34, 110
  budget planning 6, 71
  cartography 31, 76
  clinker grinding 7, 79
  cloud detection 72, 74
  cold reduction mill 156
  communications management 70
  consulting 13, 172
  dairy operation 140
  diagnosis of liver disease in small animals 10, 113
  drug reactions 114
  earthquake data analysis 9, 111
  emergency response 6, 71, 95
  exception reporting 152
  firefighting 97
  kraft mills 4, 32
  management of chemical spills 71
  management of hazardous spills 8, 95
  management of oil spill cleanup 96
  mapping 7, 77, 165
  medical diagnosis 8, 17, 43, 88, 127, 194
  medical insurance 66
  mine countermeasures 71
  mine drilling 41
  minefield detection 71
  motion understanding 4, 37, 123, 128
  music composition 10, 119
  oceanography 153
  open-pit mining 8, 91
  ore handling 41, 78
  parasite identification 8, 91
  payload tracking, capture, berthing 5, 62
  photo interpretation 6, 17, 74, 99, 166
  pitch deposits 4, 32
  power restoration 13, 168
  problem solving 63, 141, 149, 193
  robotics path planning 8, 45, 48–49, 88
  satellite communications 70
  satellite date processing 6, 13, 93, 131, 153, 165
  search and rescue 71
  sewage system control 10, 118
  ship classification 70
  temperature control 13, 91, 169
  theorem provers 118, 193
  traffic control 97
  train control 96, 151
  treatment of equine colic 10, 112
  warfare systems 7, 84, 171
  zinc leaching 7, 79
Applied AI Systems, Inc 11, 142–143, 197
Architecture 48
  blackboard 113
  computational 24
  connectionist 106, 126
  parallel 9, 64, 106, 110, 133–134
  pipeline 9
  system 148
  VLSI 106
ARIES
  See Laboratory for Advanced Research in Intelligent Educational Systems
Arithmetic, interval 9, 107
ART
  See Automated Reasoning Tool
Artificial Intelligence Research and Development Fund 51, 53, 55, 67, 74
Assembler Language 169
Associate Committee on Artificial Intelligence 14, 100, 192, 194–196
Association of Provincial Research Organizations 91
Athabasca University 99
Atlantic Canada Opportunities Agency 53
Atomic Energy of Canada Limited 41, 45
Auditing 14, 171
Australia 13, 25, 105, 112, 155, 171
Austria 113
AutoCAD 94, 160
AutoIntelligence 103
Automated inspection 8
Automated Reasoning Tool 85, 93, 113
Automatic parameter tuning 12

Automation 8, 43, 63, 92–93
  manufacturing 91
  of feature detection 70
  of industrial processes 49, 78, 147
  of knowledge acquisition 9
  of knowledge engineering 31
  of linguistic theories 125
  of marine operations 11
  of navigation 142
  of translation 67
  office 6, 66, 104
  power distribution system 168
Autonomous systems 17, 24, 62, 105–106, 194
Autonomous vehicle 8, 88, 143

Bacteriology 112
Baginski, C. 152, 195
Ballast control 8
Barron, J. 37, 130
Battle and situation knowledge 6
Bayes' theorem 10, 82, 106, 114, 123
Bayoumi, M.M. 49, 116–117
BCW Systems Ltd. 88
Belanger, P.R. 48, 148
Bell Canada 10, 119, 144
Bell-Northern Research 11, 41, 144–146
Benefits of AI 17, 19, 93, 124, 157, 175, 194
Bhattacharya, B. 124
Bilingual text generation 6
Biological system 129
Biomedicine 43, 45, 91, 123
  See also Medicine
Biotechnology 51, 94
Blackboard control structure 122
Blackboard shell 63
Blackboard systems 31
Blevis, E. 119
Blostein, D. 35, 116, 119
BMC
  See Bureau of Management Consulting
BNR
  See Bell-Northern Research
Boeing Computer Services 102
Bolc, L. 193
Bolger, L. 195
Boulet, M.-M. 195
Bowen, D.G. 154, 195
Boyle, R.W. 83
Brahan, J.W. 90
Brainmaker 77
Brassard, D. 78
Bristol Aerospace Limited 41–42
Britain 11, 25, 142
British Columbia Advanced Systems Institute 25, 41–42, 91, 167
British Columbia Hydro 41
British Columbia Research Council 91
Browse, R. 4, 35, 116–117
Buchal, R.O. 48
Bult, T. 167
Bureau of Management Consulting 7, 84
Business 3–4, 11, 14, 32, 51, 53, 137, 139, 142, 156, 173–174, 192

C 85, 94, 112–113, 139, 149
C++ 85, 154
CAD 10, 13, 31, 47, 124, 126, 156, 160
CAD/CAM 31, 93–94
CAE Electronics Ltd. 11, 41, 61, 128, 147–150
Caelli, T. 116
Calvert, T. 124
Campbell & Associates 83
Campbell, A. 83
Campbell, J. 36
Canada Centre for Mapping 6–7, 72, 76–78
Canada Centre for Mineral and Energy Technology 6, 72, 78
Canada Centre for Remote Sensing 6, 55, 72, 171
Canadian AI Capability Survey 3, 179–183
Canadian Artificial Intelligence 9, 19, 92, 100, 113, 179, 192
Canadian Astronautics Ltd. 61
Canadian Atmospheric Environment Service 95
Canadian Electrolytic Zinc 7, 79
Canadian Information Processing Society 100, 192
Canadian Institute for Advanced Research 3–5, 9, 23–26, 38–39, 41, 43, 100, 179
Canadian Marconi Company Limited 41–42
Canadian Pacific 12, 151–152
Canadian Society for Computational Studies of Intelligence 9, 92, 100, 192–193
Canadian Space Agency 5, 61–65, 88, 92, 94–95, 140
Canadian Workplace Automation Research Centre 6, 55, 66–68
CANMET
  See Canada Centre for Mineral and Energy Technology
Carleton University 87, 99
Carter, J. 122
Cartography 31, 76, 78
CASE tools 33, 123, 171
Causal network 31
CCM
  See Canada Centre for Mapping
CCRS
  See Canada Centre for Remote Sensing
Cecile, M. 113
Cement Lafarge 7, 79
Centre de recherche industrielle du Québec 8, 91
Centre de recherche informatique de Montréal 3, 31–33, 41–42, 148–149
Centre for Frontier Engineering Research 63, 94
Cercone, N. 47, 101, 124–125, 127, 193, 195
Chang, E. 91, 95, 195
CHAT
  See Conversational Hypertext Access Technology
Chiu, D.K.Y. 110
Chromatic processing 11, 129
CIAR
  See Canadian Institute for Advanced Research
Cimstation 94
Clark, B. 108
Cleary, J. 108
Coal Research Laboratories 81
Cognition 71, 126
Cognitive map construction 45
Cognitive modeling
  See Modeling
Cognitive process 126
Cognitive science 4, 25, 34, 36, 99, 116, 122, 193
Collaboration 8, 92
Collaborative research 3, 7–8, 13–14, 21, 31, 41, 53, 56, 62, 66, 76, 85, 89, 91–92, 95–96, 118, 123, 126, 156, 167, 176, 194
Collin, A.E. 50, 195
Collision avoidance 5–6, 61–62, 64–65, 97, 126

Color constancy 10, 125
Combat intelligent functions 70, 140
Communication protocols 34
Communications 6, 70, 142
  dolphins 10, 113
Communications Research Centre 6, 69
CompEngServ Ltd. 12, 153–154
Computational geometry 117, 124
Computational Intelligence 9, 100, 192–193
Computer communication 31
Computer science 17, 24, 32, 99, 194
Computer-aided design
  See CAD
Computer-Based Training Systems Ltd. 94
Concurrency control 34, 47
Conflict resolution 82
Connectionism 126
Consortia 3–4, 21, 38–42
Control 42, 48–50, 81, 162
Control structures 24
Control systems 47, 93, 96, 150
Control theory 24
Conversational Hypertext Access Technology 69, 151
Cooke, J. 122–123
CORES 12, 153
Courses 32, 104, 116
CRIM
  See Centre de recherche informatique de Montréal
Crime analysis 86
CRIQ
  See Centre de recherche industrielle du Québec
CRITTER 68
CSA
  See Canadian Space Agency
CSCSIO
  See Canadian Society for Computational Studies of Intelligence
Curriculum design 8, 63, 94, 123, 140
CWARC
  See Canadian Workplace Automation Research Centre
Cynader, M.S. 45

D&S Knowledge Systems Inc. 94
Dahl, V. 124–125, 193
Dalhousie University 25, 99
Dalkir, K. 66
Danisch, L.A. 91
Data acquisition 93
Data analysis 7, 9, 88, 94, 155–156
Data fusion 7, 70, 72, 84
Data integration 10, 55, 74, 76, 127
Data interpretation 43, 93
Data management 12, 47, 54, 63, 76, 88, 118, 153
Data presentation 63
Data processing 6, 13, 70, 93, 165, 173
Data recording 169
Data representation 10, 170
Database design 88
Database designer 125
Database management 89, 144
Database query 11–12, 32
Database research 4
Database technology 85
Dataspan Technologies Inc. 94
Davis, S. 124
De Mori, R. 33, 46, 147, 195
Decision support systems 11, 53, 71, 89, 102, 143
Deep knowledge 31, 54
Defence and Civil Institute of Environmental Medicine 62, 70
Defence Industry Productivity Program 5
Deguise, J.-C. 83
DeHaan, J. 95
Delgrande, J. 125
Dempster-Shafer theory 9–10, 70, 82, 111, 113, 123
Department of Communications 6, 55, 66, 69
Department of Fisheries and Oceans 12, 143, 153
Department of National Defence 6, 54, 67, 70–71, 86, 92, 94
Department of Supply and Services 51
Design 7, 31, 48, 50, 88, 92–93, 103, 126, 134, 144, 149, 157–158
Design systems 89
Det norske Veritas 94
Detection
  fault 62
  signal 6, 71
Diagnosis 6, 62–63, 65, 122
  equipment 157
  fault 8, 11, 63, 92–93, 95, 147–148
  medical 43, 127
  system 157
Diagnostic systems 126
  car engine 127
  electrical equipment 31
  electrical power system 63
  fault 64, 94, 151, 167
  jet engine 8, 87
  locomotive 151
  medical 8, 10, 88, 110, 112, 127
  monitoring of dams 31
  safety of earth dams 13, 159
  transformers 13, 159
  transportation 96
  veterinary screening 140
Diagnostics 64, 81
DiCiccio, V. 37, 195
Digital Equipment Corporation 12, 155–156
DIPP
  See Defence Industry Productivity Program
Discrete-event system 4, 34
Discrimination
  of consonants 35
  pattern 131
  texture based 10, 116, 131
Distribution system 13
Division of Electrical Engineering 7, 62, 85, 89
Division of Mechanical Engineering 86, 88
DND
  See Department of National Defence
Driver, J. 165
Drouin, M. 68, 195
Druce, G. 165
Dubois, A. 55, 195
Dynacon Enterprises Limited 63, 88
Dynamic manipulation 10, 118
Dynamic planning 10, 122
Dynamics 47–50, 64, 71

Earth resource satellites 13, 165
EBCO Industries Limited 41
Echidna 126
École Polytechnique de Montréal 41
Economic Regional Development Assistance 6, 66

- EDMAN 12, 153
- Education 6, 10–11, 66, 84, 88, 99, 122, 140, 171
- Effectors 24
- Efficiency and Alternative Energy Technology Branch 7, 81
- Eiffel 85
- Electromagnetic interference 8, 12, 153
- Electromagnetic Interference Expert System 12, 153
- Elio, R. 95
- Ellis, R. 35, 116–117
- ElMaraghy, H.A. 48
- ElMaraghy, W.H. 48
- EMIR
  - *See* Electromagnetic Interference Expert System
- Employment and Immigration Canada 104
- EMR
  - *See* Energy, Mines and Resources Canada
- ENERGEX 159
- Energy management 81
- Energy Research Laboratories 7, 81
- Energy, Mines and Resources Canada 6, 55, 72–83
- Engineering 31–32, 142, 157, 160
  - circuit design 11
  - computer systems 44
  - human factors 175
- Engineering applications 47
- Engineering design 42, 155
- Environment 74, 92, 172
- Environment Canada 54–55, 75, 92, 166, 171
- ERDA
  - *See* Economic Regional Development Assistance
- Error recovery 34, 106
- ETS 103
- Europe 11, 142, 155, 176
- European Economic Community 25
- Expert system consulting services 155, 171
- Expert system development 9, 11–12, 72, 80, 86, 98, 139, 155, 161
- Expert system reasoning engine 126
- Expert system shell 14, 33, 70–71, 73, 80, 85, 107, 139–140, 169, 174
- Expert systems 10, 12–13, 17, 31, 43, 61–62, 72, 91–92, 99, 102, 118, 124, 139, 144, 153, 155, 171, 173, 194
  - air drilling advisor 175
  - air traffic 12, 97, 153
  - alarm processing 156, 158, 163
  - analysis of LANDSAT data 72
  - Analyst Advisor 73
  - application of hazardous control point analysis 171
  - *See also* Applications, Advisory systems, Expert systems
  - assessment of water resources 54
  - automated data management 118
  - boreal forest experiment 75
  - building management 7
  - CAD assistant 156
  - capital acquisition 71
  - case writing 55
  - clinker grinding 7, 79
  - coal assessment 7, 81
  - collision avoidance 5, 61–62
  - computer-integrated manufacturing 103
  - configuration 94
  - conflict resolution 153
  - CORES 12
  - dam surveillance 159
  - data management 75
  - data modeling 32
  - diagnostic 32, 81, 159
  - education and training 6
  - electromagnetic hardening 88
  - electromagnetic interference 12, 153
  - emergency response advice 71
  - ENERGEX 159
  - energy exchange management 161
  - energy management 7
  - Enhanced Energy Simulation for Low Rise Buildings 81
  - environmental monitoring 75
  - equipment diagnosis 159
  - event interpretation 158
  - EXSEMBLE 160
  - fault diagnosis 64
  - financial underwriting 156
  - for design 44
  - forest inventory 55, 73–74
  - foundry technology 7, 81
  - furnace operation 7
  - haematology 139
  - hazardous spills management 95, 97
  - help desk support 14, 171, 175
  - hierarchical 7, 73
  - highway intersection design 94
  - highway location 97
  - information fusion 6, 70
  - information management 6, 66
  - inventory management 156
  - know-how management 66
  - knowledge acquisition 6
  - knowledge capture 14, 173–174
  - knowledge management 6
  - knowledge modeling 6
  - Large Building Energy Efficient Design 81
  - lube oil analyst 151
  - maintenance of CF-18 94
  - maintenance of diesel engines 12, 151
  - maintenance of rail engines 96
  - manpower scheduling 88
  - marine inspection 97
  - market analyis 139
  - medical 13, 32, 53, 88, 112, 171
  - MEDRISK 53
  - mineral processing 7, 78
  - mould cavity design 7
  - neuropsychology 139
  - nuclear power plant 13
  - off-line 79
  - oil industry 14
  - on-line 79
  - operations planning 63
  - pay and benefits 6, 66
  - personnel policy 14, 174
  - petroleum products pipeline scheduling 156
  - PIKES 74
  - planning system 64
  - plant design 54
  - plant operation 78
  - power plant system status monitor 170
  - power utility 13, 31, 156–157, 168
  - process control 4, 7–8, 41, 72, 79, 87, 124
  - process operator advisor 175
  - process planning 95
  - production of substation layouts 160
  - production rule 7, 82
  - proposal evaluation 6, 66, 171

  QUARTZ 66
  radio fault diagnosis 12, 151
  RAISON/AI 54
  real time 41, 87
  recovery of coal fines 7, 81
  resistance welding 8, 88
  resolution of interference between radio systems 97
  royalty rules 14, 174
  rule based 10, 95, 122, 170
  rule making 159
  SAGE 160
  scheduling of hot metal distribution 156
  school psychology 139
  screening 13
  SEDA 159
  selection of metal alloys 7
  SEPIDE 158
  set-up assistant for a cold reduction mill 156
  SHERI 7, 73
  sports 103
  statistical mapping advisor 77
  SUBAREX 159
  SWIFT 95, 97
  system maintenance advisor 175
  test data generation 169
  threat evaluation and weapon assignment 70
  toxicity assay 7, 84
  toxicology advisor 175
  training 74
  TRANSEPT 158
  transmission system design 158
  transportation 8, 97, 142
  user interface 175
  vessel detection and classification 6, 70
  weather forecasting 13, 37, 54, 95, 97, 156
  well cementing advisor 175
  well log analysis 8, 94
  well test interpretation 94
  zinc plant control 7
EXSEMBLE 160

Falconbridge Limited 41
Farkas, J. 66
Farrett, P. 119
Feasibility study 4–5, 32, 39, 53, 55, 62, 67, 87, 94, 139, 147, 168
Feature Extraction 165–166
Ferihill Technologies Ltd. 76
FEX 166
Financial operations 14, 39, 140, 155–156, 172–174
Finkelstein, A. 55
Finning Caterpillar 167
Fiskel, F. 53
FLIGHT
  See Flight Intelligent Troubleshooting Environment
Flight Intelligent Troubleshooting Environment 12, 147–149
Flight simulator 4, 135, 148
Fong Lochovsky, A. 110
Forestry Canada 75
Fortran 72–73, 168–169
France 25
Fuchi, K. 193
Funding 3, 21, 26, 38, 40, 51, 53, 55–56, 74, 100, 113, 119, 128
Funt, B. 124–125
Fuzzy approach 113
Fuzzy data 54

Gaines, B. 102–103, 108
GALEN 127
Gamble, K. 165
Garrett, R. 83
GCG
  See Government Consulting Group
Generalization 7, 76, 86
Geological Survey of Canada 6–7, 72, 81
Geology 44, 140
George, P.G. 165, 167, 195
Germany 25
Gershon, R. 128–129
Glasgow, J. 116–117
Gnu EMACS 107
Goldenberg, A.A. 49
GoldWorks 85, 93, 168
Goodale, M. 36
Goodenough, D.G. 55, 76, 83, 195
Government 40, 56, 100, 142, 192
  British Columbia 75
  Ontario 171
Government Consulting Group 7, 84
Government Support Programs 4, 51–55
Graham, D.P.W. 65
Graham, G. 156, 196
Grasping 48, 50, 117, 126
Greer, J. 122–123
Griffith, A. 143, 196
Grosz, B. 193
GSC
  See Geological Survey of Canada
Gupta, K. 126
Gurnsey, R. 36

H.A. Simons Limited 41
Hadley, R. 126
Haematology 11, 139
Hallett, P.E. 128–129, 131
Han, J. 124
Haptic attention 36, 116
Harauz, G. 110
Hassner, D. 167
Hatch Associates Limited 41
Havens, W. 126
Hayes, K. 195
Hazardous effluent 54
Hazardous environment 13, 44–45, 158
Hazardous spills management 8, 95
Health and Welfare Canada 53
Health care 48, 93, 171–172
Help desk 14, 151, 171, 175
Help system 123
Hewlett Packard Laboratories 92
Highway design 8, 97
Hill, E. 170, 196
Hinton, G.E. 4, 35, 46, 128–129, 135
Hirst, G. 192
Holland 25
Hollerbach, J. 50
Hostile environment 43–44
Hughes, P.C. 49
Human factors 10, 44, 50, 63, 96, 122, 175
Humphrey, G.K. 36
Hunt, W. 113
Hunter, D. 65
Hunter, I.W. 48

Hydro-Québec 12, 41–42, 157–164
Hydro-thermal scheduling 13, 168
Hymarc Ltd. 64

IBM 92, 118
Image acquisition 45
Image analysis 4, 8, 11, 37, 72–73, 91, 130, 132, 166
Image classification 74
Image databases 34
Image interpretation 45, 119, 131
Image processing 6, 9–10, 13, 45, 71, 73, 110, 122–123, 144, 165–166, 171
Image reception 13, 165
Image representation 11, 129
Image segmentation 10, 36, 116
Image understanding 17, 130
IMP Group Limited 61
Inco Limited 41
Industrial materials 8, 51
Industrial Materials Research Institute 89
Industrial Research Assistance Program 5, 8, 40, 53, 85, 88
Industry 3, 11, 14, 40, 45, 49, 54, 56, 100, 137, 192, 194
  aerospace 11, 39, 48, 92–93, 139
  agriculture 11, 49, 93, 139
  Canadian 4, 51–55, 88
  computing 9
  construction 8, 49, 89, 92, 140
  electrical 157
  energy 78
  entertainment 93
  forestry 11, 48–49, 55, 139, 167
  hydroelectric 4, 13, 31
  manufacturing 17, 32, 43, 45, 49, 52–53, 87, 92–93
  metallurgical 32
  metals 78
  minerals 78, 140
  mining 8, 11, 39, 41, 48–49, 54–55, 91, 139
  oil 8, 14, 39, 48, 93–94, 156
  power 39, 93, 147
  process control 32, 148
  process manufacturing 155
  processing 7, 39, 52–53, 78
  resource 7, 43, 55, 74, 87, 93
  service 32, 43, 49, 52–53
  transportation 8, 11, 43, 92, 94, 96, 140, 142, 153
Industry, Science and Technology Canada 5, 40, 51, 53, 55, 73–74, 179
Inference mechanisms 44, 93, 111, 148, 173
Inferencing 111, 174, 176
Information system 14, 44, 46, 56, 122, 144, 172
  executive 172
  geographic 55, 82
  large scale 123
  manufacturing 34
  medical 112
  on-orbit 62
Information technology 34, 47, 51, 89, 171
Information Technology Research Centre 4, 25, 34–37, 119
Injection molding 156
Insight 2+ 153
Inspection
  industrial 9, 34, 110
  visual 91, 110
Institute for Information Technology 62, 85
Institute for Research in Construction 8, 89, 92
Institute for Robotics and Intelligent Systems 4–5, 9, 25, 40, 57, 100, 122–123
Institute of Aerospace Research 62
Insurance policy screening 10
intelCAD 126
Intelligent tutor 71, 88, 97–98
Intera-Kenting Technologies 76
Interactive Tactical Environment Management System 12, 147, 149
Interdisciplinary research 9–10, 14, 24–25, 35, 56, 67, 94, 102, 110, 116, 119, 123–124, 128, 163
Interface
  adaptive 46
  DECwindows 73
  for disabled people 105
  graphic 42, 44, 47, 149
  human–computer 36, 70, 102, 106
  intelligent 31
  interactive 107
  man–machine 31, 124
  natural language 8, 11, 44, 47, 86, 97, 125, 144
  quasi-intelligent 71
  query 125
  robot 49
  speech 6, 17, 44, 47, 69
  text 44, 47
  user 12, 44, 47, 63, 76, 87, 103–104, 113, 142, 148, 153, 162, 175
  user-friendly 34
  vocal 32
International Submarine Engineering Ltd. 64, 95
IRAP
  See Industrial Research Assistance Program
IRIS
  See Institute for Robotics and Intelligent Systems
Isabelle, P. 55, 67
ISE Research Limited 63
ISIM 95
ISTC
  See Industry, Science and Technology Canada
ITEMS
  See Interactive Tactical Environment Management System
ITRC
  See Information Technology Research Centre

Jacobson, Z. 84, 196
Japan 11, 25, 142, 155, 157
Jenkin, M.R.M. 128–130
Jenkins, M. 117
Jennings, R. 126
Jepson, A.D. 128–131
Johnson Matthey Electronics 124
Jones, J. 126
Jones, M. 122

Katz, L. 108
KEE
  See Knowledge Engineering Environment
Kidd Creek Mines 7, 79
KITTEN 103, 106
Knowledge acquisition 4, 6, 9, 11, 23, 31, 34, 61, 66–67, 70, 86–87, 99, 102–104, 111, 139–140, 153, 174–175
Knowledge base
  large 31–32, 34
  object oriented 103

Knowledge capture 14, 174
Knowledge dissemination 4
Knowledge engineer 12, 67, 139, 155, 162
Knowledge engineering 4, 9, 12, 31, 70, 80, 102, 104, 122, 155, 171
Knowledge Engineering Environment 12, 79, 93, 145, 153, 171
Knowledge House Publishing Ltd. 88
Knowledge management 6, 66
Knowledge manipulation 23
Knowledge modeling 6, 66, 119
Knowledge organization 34
Knowledge processes 103
Knowledge representation 10, 18, 23–24, 26, 34, 44, 47, 67, 93, 95, 99, 113, 122–127, 144, 148, 162, 193
Knowledge science 9
Knowledge synthesis 34
Knowledge worker 6, 66, 176
Knowledge-base construction 11
Knowledge-base design of information systems 44
Knowledge-base development 67, 78, 139
Knowledge-base generation 86
Knowledge-base interface 44
Knowledge-base management 34, 44, 46
Knowledge-based modeling 10
Knowledge-based system development 89, 118
Knowledge-based systems 5, 7–10, 12–13, 17, 31–32, 43, 46, 85, 100, 102, 116–117, 148, 165, 171
  See also Advisory systems, Applications, Expert systems
  aircraft maintenance 118
  automatic software testing 118
  character recognition 94
  chess tutor 122
  contract evaluation 142
  decision making 11
  depreciation 118
  drug dispensing 9, 104
  engine monitoring 86
  fault diagnosis 86, 167
  fault prediction 13
  fuzzy sewage and system control 118
  ground support 7
  image extraction 13, 166
  insurance policy screening 118
  mining 81
  nuclear power plant 13, 169
  office 66
  power utility 147
  process engineering 118
  sports 9, 103
  tax 171
  thermal system control 119
  tutoring 140
  urban transit 142
  vehicle dispatcher's assistant 142
  weather forecasting 165
  well logging 118
KnowledgeCraft 93
Knowledgeworks Research Ltd. 88
Kolers, P. 128
Konkin, D. 165
Krause, M. 165, 167
KSSO 103
Kurusek, A. 132
Kwan, H. 128–129, 134
Kwok, P. 108

Laboratory for Advanced Research in Intelligent Educational Systems 10, 122
Laboratory for Intelligent Systems 7, 85, 88
Lac Minerals Limited 41
Laguitton, D. 80
Lam, D. 54
LANDSAT 7, 13, 72, 165
Lane, A. 83
Laugesen, E. 173, 196
Laval University 87–88, 99
Lavalin 7, 79
Lawrence, P.D. 49
Learning 8, 23, 99, 113, 134, 193–194
  by example 86
  collaborative 93
  computer-assisted 85
  computer-managed 93–94
  connectionist 46
  explanation-based 105, 123
  incentive 71
  inductive 34, 111
  See also Machine learning
Learning procedures 4, 35, 46, 135
Learning systems 8, 10, 88, 92, 106, 124
Lederman, S. 35, 44, 116–117
Lemire, A. 164, 196
Leung, J. 80
Levine, M.D. 45
Lewis, R.T. 98, 196
Li, Z.-N. 127
Linguistics, computational 10, 124, 127
LISP 12, 79, 85, 95, 107, 113–114, 135, 139, 145, 153, 166, 168, 174
Lockheed Software Technology Center 102
Logic 162
  deductive 23
  first order 125
  fuzzy 8, 10, 65, 91, 106, 123
  probabilistic 9, 111
  symbolic 10, 23, 119
Logic depiction 131
Logic grammars 127
Logic programming 9, 18, 44, 47, 102, 106–107, 118, 124–127, 193
Lopez, A. 113
Lumb, C. 95

MacDonald Dettwiler and Associates Ltd. 13, 41, 61, 92, 95, 165–167
MacDonald, B. 108
Machine intelligence 8–9, 93, 110
Machine learning 6, 8–9, 64–65, 70, 74, 86, 93, 102, 104–105, 113, 126, 139–140, 156
Mackworth, A.K. 44, 131, 193
MacLean, S. 65, 196
MacMillan Bloedel Ltd. 13, 140, 167
MacNabb, G.M. 42, 195
Maintenance 92
  aircraft 10, 118
  Automated Defence Digital Network 70
  CF-18 8, 94
  diesel engine 12
  drilling platform 63
  electrical equipment 31
  electro-mechanical system 165

  jet engine 8, 87
  knowledge-based system 172
  large information system 46
  locomotive 96, 151
  logging trucks 13
  marine 142
  power lines 32
  railway track 12, 151
  space station 62
  telephone switch 11
  towed submersible cable faring 63
  transmission line 4
  vehicle 9
Manalta Coal Limited 41
Manipulators 5, 24, 44, 49–50, 61–62, 64, 93, 117, 126
Manning, E.G. 167, 195
Manolis, R. 164, 195–196
Manufacturing 155
  computer-integrated 93, 103
Manufacturing Research Corporation of Ontario 118–119
Manufacturing systems 56, 93
  supervisory 86
Marine environment 12, 153
Marine inspection 97
Marine operations 11, 142
Market research 11, 53, 139–140, 155–156
Matching 131
  image 34
  shape 36, 43, 131
Mathiel, R. 91
Mayman, S.A. 196
McCalla, G. 101, 122–123, 193, 195–196
McCarney Technologies 127
McGill University 3, 25, 41–42, 45, 57, 99–100, 123, 147–149
McLeish, M. 110–114
McMaster University 48, 87
Medical Research Council 57
Medicine 8, 11, 44, 139
  See also Biomedicine
MEDRISK 53
MEDUSA 61
Meetings 25, 100
Memorial University 99
Mendelzon, A.O. 46
Mercer, R. 37
Metals Technology Laboratories 7, 81
Métavision Inc. 86
Meteorology 4, 13, 32, 37, 165
Microelectronics and Systems Development Program 5, 51
Microlight Corporation 126
Microtel Pacific Research Limited 41
Milios, E. 128–131
Mineral processing 7, 39, 78–79
Mineral resource assessment 7, 82
Mineral Science Laboratories 7, 78
Mining Research Laboratories 7, 80
Mobile Servicing Centre 5, 61
Mobile Servicing System 5, 10, 50, 61–63, 65, 95, 118, 140, 167
Modeling 92–93, 102, 179
  brain 71
  cognitive 10, 122, 126
  data 32, 106
  decision processes 149
  dynamics 49
  human hand movement 116
  See also Knowledge modeling
  knowledge-base 119
  neural 45
  qualitative 92, 94
  river 13, 168
  rule network 34
  shape 131
  speech 46
  text 105
Modi, V.J. 50
Monitoring
  machinery health 86
  real time 42
  system health 62
Monitoring systems 93
  ballast control system 94
  emergency operating procedures 169
  environmental 74
  power 159
  systems status 169
Moore, T. 116
Moraglia, G. 128–129, 134–135
Motion analysis 10, 37, 99, 122
Motion extraction 129
Motion planning 10, 48, 124, 126
Motion representation 35
Motion sensing 11, 129
Motion understanding 4, 123, 128
MPB Technologies Inc. 41–42, 64, 87
MRCO
  See Manufacturing Research Corporation of Ontario
MSC
  See Mobile Servicing Centre
MSDP
  Microelectronics and Systems Development Program 4
MSL
  See Mineral Science Laboratories
MSS
  See Mobile Servicing System
Mulder, J. 195
Music 10, 116
Muzard, J. 66
Mylopoulos, J. 4, 34, 46, 113, 128, 193

National Aeronautical Establishment 62
National Research Council 5, 7, 14, 40–41, 51, 53, 62, 85–90, 92, 96, 100, 119, 179, 192–196
Natural language processing 17, 68, 97, 144, 194
Natural language research 32, 34, 69, 127, 142, 144, 173
Natural language system 10, 113
Natural language text 55
Natural language translation 67
Natural language understanding 10–11, 95, 99, 122, 124–125, 193
Natural Sciences and Engineering Research Council 5, 56–57, 113, 119, 122
Navigation 44, 142
  air systems 8, 96
  robot 41, 43, 45, 130
Neate, J. 54
Nested Interactive Array Language 10, 63, 118
Networks of Centres of Excellence Program 4, 25, 41, 43, 47, 57, 100, 113, 122–123
Neufeld, E. 122–123
Neural net 4, 7–9, 84, 91, 96–98, 106

Neural network 6–7, 11–12, 35, 46, 50, 61, 64, 69, 71, 82–84, 96, 126, 142, 153, 156
Neuronal network 134
Neuropsychology 11, 139
New Brunswick Power Commission 13, 41, 168–170
New Brunswick Research and Productivity Council 8, 91
Nexpert 79, 85, 93, 103
Nial
  See Nested Interactive Array Language
Nickerson, B. 78
Noranda Inc. 41, 87
Norrie, D. 108
Northern Telecom 11, 92, 95, 144
NRC
  See National Research Council
NSERC
  See Natural Sciences and Engineering Research Council
Nuclear power plant 11, 13, 147, 169

Oerlikon Aerospace Limited 41, 64
OITMIS 62
Ontario College of Veterinary Medicine 112
Ontario Hydro 41, 45, 147
Ontario Premier's Council Technology Fund 4, 34
Optech Inc. 64
Optimization
  mathematical 12, 148
  numerical 31, 149
  query 34
Optimization problems 71, 132
Oracle 77, 112–113
Organization principles 45
ORTECH International 41
Ortony, A. 193

Panabode Corporation 126
Pang, G.K.H. 4, 34
Parallel computing 31, 133
Parallel networks 35, 135
Parallel processing 4, 36, 50
Parasitology 112
Pascoe, P. 113
Path planning 8, 45, 48–49, 62, 64–65, 88
Pattern analysis 8–9, 93, 110
Pattern matching 173
Pattern recognition 9, 70, 93, 110, 135, 156, 165
Patterson, B.K. 170, 196
Payload operations 5, 61–62
Peacocke, R. 144, 195
Pension policy 14, 174
Perception 10, 23–24, 44, 116
  auditory 43
  color 124–125
  computational 5, 43–44, 47, 116
  motion 36
  multisensor 45
  object 117
  shape 36
  speech 46
  tactile 35, 43–44
  visual 44
Perrault, R. 193
Personnel policy 14, 174
Petro-Canada Resources 41
Phillips, D. 69, 196
Photo interpretation 6, 17, 74, 99, 166
Picot, J. 55
PIKES 74
Planning 6, 48, 62–65, 95, 105, 157, 179, 193
  allocation of resources 167
  budget 6, 71
  business 4
  See also Motion planning
  operational 9, 88
  production 89
  task 48
  trajectory 6
Planning system 5, 12, 61, 63, 83, 89, 93, 153
Poiker, T. 124
Point Lepreau Nuclear Generating Station 13, 169–170
Poplog 85
Popowich, F. 127
Posluns, H. 54
Power distribution system 168
PRECARN
  See PRECompetitive Applied Research Network
Precipitation
  See Meteorology
Precipitation Forecasting Meteorologist's Assistant 165
PRECompetitive Applied Research Network 4–5, 11, 25, 38–44, 57, 87, 92, 113, 147, 170, 176, 179
Prediction 11, 102
  failure 63, 81, 87
  fault 13, 167
PREFORMA
  See Precipitation Forecasting Meteorologist's Assistant
Price Waterhouse 13, 76, 88, 171–172, 179
Process control 4, 8, 10–11, 32, 41, 78–79, 87, 93, 124, 147–148
Process engineering 118
Process planning 8, 17, 86, 92, 95, 194
Programming
  by example 9, 102, 104
  object oriented 12, 93, 114
Programming environments 10, 99, 107, 122
Programming language 5, 9, 12, 47, 61, 118–119, 153
Prolog 13, 47, 68, 72–73, 83, 85, 93, 107, 124, 145, 166, 168, 174
Prototyping 9, 46, 78, 80, 87, 102, 106, 144–145, 151
Provincial Research Organizations 8, 91–95
Pylyshyn, Z. 4, 26, 36–37, 193, 195

Q'HAND
  See Queen's Haptic Attention Device
Q'Nial 112
  See also Nested Interactive Array Language
QUARTZ 66
Queen's Haptic Attention Device 116
Queen's University 3–4, 10, 34–36, 41, 44, 49, 99, 116–121

Rach, C. 91
Radar 6, 71, 97
  synthetic aperture 13, 70, 165
  weather 4, 37
Radiology 112
RAISON/AI 54
Realtime processing 8, 10, 13–14, 48, 62–65, 87, 93, 124, 143, 147–148, 153, 158–160, 168
Reasoning 24, 26, 32, 44, 47, 95, 124, 126, 131, 194
  assumption-based 179
  automated 47
  case-based 64, 93

  common sense 47, 67, 123
  constraint-based 92
  default 112
  heuristic 54
  metaphorical 113
  nonmonotonic 9, 111
  probabilistic 10, 122
  procedural 5, 61
  qualitative 10, 12, 31, 122, 148–149
  syllogistic 23
  temporal 92, 95, 166
Recognition 194
  character 8, 91, 94
  granularity-based 122
  image 156
  object 6, 34, 36, 44, 70, 116–117, 132
  See Pattern recognition
  phenome 84
  shape 4, 10, 35, 123, 135
  signal 7, 84
  speech 4, 6–7, 12, 31–32, 35, 69, 84, 124, 147, 149–150
  surface 36
  tactile 10
  voice 9, 17, 97–98, 149
Regan, D.M. 129, 135
Reimer, W. 176, 196
Reiter, R. 47, 128–129, 131, 193
Remote sensing 7, 13, 43, 55, 72, 165–166
Remote Sensing Shell 7, 72
Rendell, L. 193
Repair
  diesel engine 12
  electrical equipment 31
  rail lines 151
  transmission line 4
Representation 95, 131
  See Data representation
  See Image representation
  See Knowledge representation
  object 7, 34, 43, 76
  of uncertainty 105
  shape 129
RESHELL
  See Remote Sensing Shell
Richardson, D.E. 77, 83
Risk assessment 14, 53, 88, 171
Robotics 3–5, 8, 10, 12, 24, 26, 38–39, 41, 43–45, 48–49, 70, 91–94, 99–100, 116–117, 153, 170, 179
  autonomous 5, 63
  sensor-based 8, 85
  teleoperated 5, 63, 117
Robotics Systems International Ltd. 64
Robots 6, 9–10, 17, 24, 48, 104, 117, 126, 130, 194
  autonomous 35–36, 44
  dexterous 49, 61
  drug dispensing 104
  ground-based 50
  high performance 48
  intelligent 5, 32, 41, 43–44, 47–48, 116, 126
  mobile 41
Roesler, R. 167
Rosenberg, R. 167
Royal Bank of Canada 14, 173
Rule-based system 12, 14, 95, 97, 122, 151, 153, 170
Safety 140
  air 8, 96
  building 81
  mine 81
  of earth dams 13, 159
  of medical devices 53
  rail 12, 96, 151
SAGE 160
Sandwall, E. 193
Saskatchewan Research Council 91
Sauder, B. 167
SCENT
  See Student Computing Environment
Scheduling 150, 156
  budget 71
  constraint-based 89
  fermentation plant 8, 94
  fleet 92
  hydro-thermal 13, 168
  manpower 88–89
  operations 17, 194
  order 156
  pipeline 156
  public transport 97
  rail cars 12
  repairs 151
  track maintenance 12
  training courses 67
Scheme 107
Schubert, L. 193
Schwab, D. 167
Scrimgeour, J. 195
Security, air 8, 96
SED Systems Inc. 61, 92
SEDA 159
Segmentation 10, 36, 74, 122, 129, 131
Segregation 131
Sekely, G.F. 195
Selective attention 4, 34
Seminars 7, 14, 32, 57, 84, 92, 96–97, 128, 171
Sensing 43
Sensor Adaptive Machines Inc. 64
Sensor fusion 43, 63, 93–94, 179
Sensors 17, 24, 45, 49, 70, 93, 117
Sensory integration 43, 45, 49, 93, 126–127
Sensory processing 45
SEPIDE 158
Severe Weather Intelligent Forecast Terminal 95, 97, 165
Shape analysis 10, 122
Shape theories 131
Shaw, M. 108
Shell Canada 14, 41, 147, 174–176
SHERI
  See System of Hierarchical Experts for Resource Inventories
Siekierska, E.M. 77
Signal processing 93
Simon Fraser University 10, 42, 99, 124–127, 167
Simulated annealing 71
Simulation 11, 48–49, 65, 78, 84, 88, 92–93, 95, 132, 147, 158, 179
  aerodynamic 12, 148
  air traffic control 147, 149
  distributed 107
  energy 81
  flight 32

  naval warfare 7
  sky-train 89
  tactical mission 12, 149
  work gang 151
Sloman, A. 193
Smalltalk 77, 85, 107
Smartcam 94
Smith, B.J. 141, 195–196
Social impact 9, 38, 43, 89, 100, 102, 157, 173
Social Sciences and Humanities Research Council 57
Software 9, 11
Software engineering 11, 31–32, 80, 144
Software reusability 46
Softwords Division of Press Porcepic 63, 88
Sonar 7, 70–71, 84
Sorenson, P. 123
Space Station 5, 8, 50, 61–62, 95, 165
Space Station Remote Manipulator System 5, 61
Spar Aerospace Limited 41–42, 61–62
SPDM
  See Special Purpose Dextrous Manipulator
Special Purpose Dextrous Manipulator 5, 61, 63
Spectrum Engineering Corporation Limited 63, 88, 118
Speech 4, 6–7, 12, 17, 31–32, 35, 44, 46, 69, 71, 84, 99, 135, 147, 149–150, 193
  See also Interface
  See also Modeling
  See also Perception
Speech processing 144
Speech recognition
  See also Recognition
SPOT 13, 73, 165
Sridharan, N.S. 193
SSRMS
  See Space Station Remote Manipulator System
Standards 160
STAR 165, 167
Statistics Canada 6, 68
STEAR
  See Strategic Technologies for Enhanced Automation and Robotics
STELCO Inc. 41
Stepanenko, Y. 47
Stirtzinger, T. 113
STP
  See Strategic Technologies Program
Strategic Technologies for Enhanced Automation and Robotics 5, 61, 88, 92, 94–95, 118–119, 140
Strategic Technologies Program 5, 40, 51
Student Computing Environment 122
SUBAREX 159
Substitution of data 7
Surgenor, B. 117, 119
Surveys 7, 11, 83–84, 139, 142, 148, 150, 179–183
Sweden 25
SWIFT
  See Severe Weather Intelligent Forecast Terminal
Swinkels, F. 167
System design 11, 36, 43–44, 46–47, 122–123, 131
System integration 103
System management 14, 172
System of Hierarchical Experts for Resource Inventories 7, 55, 73, 76
System organization 9, 102, 105–106
System theory 9, 102, 105
Systèmes XA Inc. 88

Tax practise 14, 171
Tax rules 14, 174
Taylor, V.K. 71, 196
Technology Outreach Program 5, 53
Technology transfer 48, 53, 56, 92, 96, 147
Telecommunications 6, 11, 32, 39, 66, 93, 144, 155
Teleoperation 49–50, 117
Telephone switching 8, 11, 94
Telepresence Research Ltd. 64
Telerobotics 4, 24, 36, 41, 43–45, 47
Testing
  of expert systems 4, 31–32, 144
  of software 10, 118
TeX 107
Text compression 9, 105
Text generation 32
Text processing 142
Text retrieval 156
The University of Western Ontario 3–4, 25–26, 34, 36–37, 48, 99
Thomas, G. 195–196
Thomson CSF Systems Canada Ltd. 64
Threat evaluation and weapon assignment 6, 70
TOP
  See Technology Outreach Program
Topographic mapping 7, 77
Toxicity assays 7
Toxicology 14, 175
Tracking 44
  object 4, 36
  storm 4, 37
Train dispatching system 12, 151
Training 6, 10–11, 55, 66–67, 81, 89, 104, 122, 140, 149–150, 155, 157–158, 172
TransAlta Utilities Corporation 41
TransCanada Pipelines 41
TRANSEPT 158
Translation
  automatic 10, 125
  language 125
  machine 6, 17, 67, 124, 194
  machine-aided 55, 67
Translator's workstation 6, 67
Transport Canada 8, 54, 96–98, 142, 150, 153
Tremblay, P. 122–123
Trent University 99
Tsotsos, J.K. 4, 34, 45, 128–130, 132–134, 193
Tutorials 84
Tutoring 6, 9, 11, 63, 65, 67, 70, 92, 99, 122, 124, 139–140

U.S. Office of Naval Research 119
Uncertainty management 9–10, 105–106, 110–111, 117, 122–123, 127
United Kingdom 11, 13, 144, 157, 171
United States 11, 13, 25, 83, 142, 144, 155, 157, 171
Université de Montréal 67, 99
Université du Québec 99
University of Alberta 3, 99
University of British Columbia 3, 25, 41, 99–100, 131, 167
University of Calgary 9, 93, 99, 102–109
University of Guelph 9, 99, 110–115

University of Leeds 88
University of Manitoba 99
University of New Brunswick 13, 41, 78, 87, 170
University of Ottawa 41, 76, 87, 99
University of Regina 99
University of Saskatchewan 10, 93, 99, 122–123
University of Toronto 3–4, 10, 25, 34–35, 41, 49, 76, 99–100, 118, 128–136, 147, 192
University of Victoria 140, 167
University of Waterloo 4, 34, 41, 62, 87, 99–100, 118

Vadeko International Inc. 64
Validation 65
  of data 7, 79
  of expert systems 4, 31–32
vanEnden, M.H. 47
Vehicle dispatching 11, 96, 142
Verification 32, 148
Veterinary medicine 10, 140
Vickers, J. 108
Virtek Ltd. 64
Vision 10, 26, 93, 125
  active 43, 45
  attentive 135
  autonomous systems 17
  biological 4, 10, 34, 45, 128–136
  computational 4, 10, 34–36, 43, 70, 99–100, 124, 127–136, 193
  computer 6, 37, 45, 70, 72, 74, 110, 119, 122–123, 130, 134
  human 35–36, 123
  machine 8, 17, 34, 37, 61–62, 91, 110, 126
  model based 132
  preattentive 135
  robot 45, 127
  stereoscopic 70, 117
Vision systems 6, 12, 44, 62, 64–65, 129, 132, 153
  active 35
  industrial 9, 110
  robot 10
  space 5
Voice interactive technology 11
Voice recognition
  See Recognition
Voice synthesis 17

W.G. Hutchison & Company Limited 41
Walden, J. 196
Weather forecasting 8–9, 13, 37, 54, 95, 131, 156, 165–166
Webber, B.L. 193
Welding 8, 88
Well log 10, 94, 118
Western Diversification Program 53
Wilkins, D. 193
Winne, P. 124
Wiring diagrams 10, 119
Witten, I. 102, 108
Wong, A. 4, 34
Wonham, W.M. 4, 34
Woodham, R. 193
Workshops 9, 25, 32, 40, 57, 92, 97, 100, 103, 192

Xerox Research Centre of Canada 41

Yang, H. 122–123
Yao, P. 113
Yee, B. 166
York University 10, 26, 41, 128

Zaky, S. 134
Zelek, J. 166
Zucker, S. 45